



Services Interfaces Library for Routing Devices



Published: 2015-03-25

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Services Interfaces Library for Routing Devices
Copyright © 2015, Juniper Networks, Inc.
All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About the Documentation	lvii
	Documentation and Release Notes	lvii
	Supported Platforms	lvii
	Using the Examples in This Manual	lvii
	Merging a Full Example	lviii
	Merging a Snippet	lviii
	Documentation Conventions	lix
	Documentation Feedback	lxi
	Requesting Technical Support	lxi
	Self-Help Online Tools and Resources	lxi
	Opening a Case with JTAC	lxii
Part 1	Services Interfaces Overview	
Chapter 1	Overview	3
	Understanding Services PICs	3
	Adaptive services and Multiservices PICs	3
	Encryption Services (ES) PIC	4
	Multilink Services and Link Services PICs	4
	Monitoring Services PICs	4
	Tunnel Services PIC	5
	Multiservices MIC and Multiservices MPC	5
	Multiservices MIC and Multiservices MPC (MS-MIC and MS-MPC) Overview	5
	Supported Platforms	7
Chapter 2	Configuration Overview	9
	Services Interface Naming Overview	9
	Enabling Service Packages	11
	Layer 2 Service Package Capabilities and Interfaces	14
	Services Configuration Procedure	15
	Example: Service Interfaces Configuration	16
	Configuring Default Timeout Settings for Services Interfaces	19
	Configuring System Logging for Services Interfaces	20
Part 2	Adaptive Services Overview	
Chapter 3	Adaptive Services Overview	25
	Adaptive Services Overview	25
	Packet Flow Through the Adaptive Services or Multiservices PIC	27

Chapter 4	Adaptive Services Configuration Overview	29
	Understanding Service Sets	29
	Configuring Service Sets to be Applied to Services Interfaces	31
	Configuring Interface Service Sets	31
	Configuring Next-Hop Service Sets	33
	Determining Traffic Direction	34
	Interface Style Service Sets	34
	Next-Hop Style Service Sets	35
	Configuring Service Rules	36
	Configuring Service Set Limitations	37
	Enabling Services PICs to Accept Multicast Traffic	38
	Applying Filters and Services to Interfaces	38
	Configuring Service Filters	39
	Example: Configuring Service Sets	41
	Configuring AS or Multiservices PIC Redundancy	41
	Examples: Configuring Services Interfaces	44
	Configuring the Address and Domain for Services Interfaces	45
	Configuring System Logging for Service Sets	47
	Tracing Services PIC Operations	48
	Configuring the Adaptive Services Log Filename	49
	Configuring the Number and Size of Adaptive Services Log Files	49
	Configuring Access to the Log File	50
	Configuring a Regular Expression for Lines to Be Logged	50
	Configuring the Trace Operations	50
 Part 3	 Translating IP Addresses Using NAT	
Chapter 5	NAT Overview	55
	Junos Address Aware Network Addressing Overview	55
	NAT Concept and Facilities Overview	56
	IPv4-to-IPv4 Basic NAT	57
	Basic NAT	57
	NAPT	57
	Static Destination NAT	57
	Twice NAT	57
	IPv6 NAT	58
	Application-Level Gateway (ALG) Support	58
	NAT-PT with DNS ALG	58
	Dynamic NAT	59
	Stateful NAT64	59
	Dual-Stack Lite	59
	Junos Address Aware Network Addressing Line Card Support	60
	Junos OS Carrier-Grade NAT Implementation Overview	60
	Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card	61

Chapter 6	NAT Configuration Overview	65
	Network Address Translation Configuration Overview	65
	Configuring Source and Destination Addresses Network Address Translation	
	Overview	66
	Configuring Pools of Addresses and Ports for Network Address Translation	
	Overview	67
	Configuring NAT Pools	67
	Preserve Range and Preserve Parity	68
	Specifying Destination and Source Prefixes without Configuring a Pool	68
	Network Address Translation Rules Overview	69
	Configuring Match Direction for NAT Rules	70
	Configuring Match Conditions in NAT Rules	70
	Configuring Actions in NAT Rules	71
	Configuring Translation Types	73
	Configuring Service Sets for Network Address Translation	75
	Carrier-Grade NAT Implementation: Best Practices	76
	Use APP and Round-Robin Address-Allocation	77
	Do Not Use EIM with SIP	77
	Do Not Use EIM with HTTP, DNS, or When Not Needed	78
	Define PBA Blocks Based on User Profiles	79
	Do Not Change the PBA Configuration on Running Systems	79
	Do Not Allocate Excessively Large NAT Pools	80
	Configure the System Log for PBA Only When Needed	81
	Use Redundant Service PIC (RSP) Interfaces for Failover	83
	Contain the Effects of Missing IP Fragments	84
	Do Not Use Configurations Prone to Routing Loops	84
Chapter 7	Avoiding IPv4 Exhaustion Using Junos Address Aware Network Addressing and Stateful NAT64	87
	Sample IPv6 Transition Scenarios	87
	Example 1: IPv4 Depletion with a Non-IPv6 Access Network	87
	Example 2: IPv4 Depletion with an IPv6 Access Network	88
	Example 3: IPv4 Depletion for Mobile Networks	89
	Configuring Stateful NAT64	89
Chapter 8	Hiding Private Networks Using Static Source NAT	93
	Configuring Static Source Translation in IPv4 Networks	93
	Configuring the NAT Pool and Rule	93
	Configuring the Service Set for NAT	95
	Configuring Trace Options	97
	Sample Configuration - Static Source NAT Using a Static Pool With An Address Prefix And An Address Range	98
	Sample Configuration - Static Source Nat for One-To-One Mapping Between a Private Subnet and a Public Subnet	98
	Configuring Static Source Translation in IPv6 Networks	99
	Configuring the NAT Pool and Rule	100
	Configuring the Service Set for NAT	101

	Configuring Trace Options	102
	Example: Configuring Basic NAT44	103
	Example: Configuring NAT for Multicast Traffic	105
	Rendezvous Point Configuration	105
	Router 1 Configuration	108
Chapter 9	Making Private Servers Available Using Static Destination NAT	111
	Configuring Static Destination Address Translation in IPv4 Networks	111
Chapter 10	Allowing Components of a Private Network to Share a Single Address Using NAPT	117
	Configuring Address Pools for Network Address Port Translation (NAPT)	
	Overview	117
	Round-Robin Allocation for NAPT	118
	Sequential Allocation for NAPT	118
	Preserve Parity and Preserve Range for NAPT	119
	Address Pooling and Endpoint Independent Mapping for NAPT	119
	Address Pooling	119
	Endpoint Independent Mapping and Endpoint Independent Filtering	120
	Port Block Allocation for NAPT	121
	Secured Port Block Allocation for NAPT	121
	Interim Logging for Port Block Allocation	122
	Deterministic Port Block Allocation for NAPT	122
	Understanding Deterministic Port Block Allocation Algorithms	122
	Deterministic Port Block Allocation Algorithm Usage	123
	Deterministic Port Block Allocation Restrictions	125
	Comparison of NAPT Implementation Methods	126
	Configuring Dynamic Source Address and Port Translation in IPv4 Networks	127
	Configuring Dynamic Source Address and Port Translation for IPv6 Networks	131
	Example: Configuring NAT with Port Translation	133
	Example: NAPT Configuration for the MS-MPC	134
	Example: Dynamic Source NAT as a Next-Hop Service	138
Chapter 11	Securing Traffic Using NAT-PT and ALGs	141
	ALGs Available by Default for Junos OS Address Aware NAT	141
	Configuring Protocol Translation Between IPv6 and IPv4 Networks -	
	NAT-PT	143
	Configuring the DNS ALG Application	144
	Configuring the NAT Pool and NAT Rule	144
	Configuring the Service Set for NAT	147
	Configuring Trace Options	148
	Example: Configuring NAT-PT	150
Chapter 12	Reducing Traffic and Bandwidth Requirements Using Port Control Protocol	165
	Port Control Protocol Overview	165
	Configuring Port Control Protocol	167
	Configuring PCP Server Options	167
	Configuring a PCP Rule	168
	Configuring a Service Set to Apply PCP	169

	SYSLOG Message Configuration	169
	Example: Configuring Port Control Protocol with NAPT44	170
Chapter 13	Automatically Assigning Ports Using Port Block Allocation	177
	Secured Port Block Allocation for NAPT	177
	Interim Logging for Port Block Allocation	177
	Configuring Secured Port Block Allocation	178
	Configuring Deterministic Port Block Allocation	180
Chapter 14	Connecting Specific Ports and Addresses Using Port Forwarding	183
	Configuring Port Forwarding for Static Destination Address Translation	183
	Configuring Port Forwarding Without Destination Address Translation	186
	Example: Configuring Port Forwarding with Twice NAT	187
Chapter 15	Allocating a Few Public Addresses to Many Private Hosts Using Dynamic NAT	191
	Configuring Dynamic Address-Only Source Translation in IPv4 Networks	191
	Example: Dynamic Source NAT as a Next-Hop Service	195
	Example: Assigning Addresses from a Dynamic Pool for Static Use	197
Chapter 16	Achieving Line-Rate, Low-Latency Translations Using Inline NAT	199
	Inline Network Address Translation Overview for MPC Types 1, 2, and 3	199
	Example: Configuring Inline Network Address Translation - Interface-Service Service Set	201
Chapter 17	Monitoring NAT	209
	Configuring NAT Session Logs	209
	Monitoring NAT Pool Usage	210
Part 4	Transitioning to IPv6 Using Softwires	
Chapter 18	Softwires Overview	213
	Tunneling Services for IPv4-to-IPv6 Transition Overview	213
	6to4 Overview	213
	Basic 6to4	214
	6to4 Anycast	214
	6to4 Provider-Managed Tunnels	215
	DS-Lite Softwires—IPv4 over IPv6	215
	6rd Softwires—IPv6 over IPv4	216
Chapter 19	Softwires Configuration Overview	219
	Configuring Softwire Rules	219
	Configuring Service Sets for Softwire	220
Chapter 20	Transitioning to IPv6 Using 6to4 Softwires	223
	Configuring a 6to4 Provider-Managed Tunnel	223

Chapter 21	Transitioning to IPv6 Using DS-Lite Softwires	227
	Configuring a DS-Lite Softwire Concentrator	227
	Example: Basic DS-Lite Configuration	228
	Example: Configuring DS-Lite and 6rd in the Same Service Set	234
	Protecting CGN Devices Against Denial of Service (DOS) Attacks	241
	Mapping Refresh Behavior	241
	EIF Inbound Flow Limit	242
	DS-Lite Subnet Limitation	242
	DS-Lite Per Subnet Limitation Overview	242
	Configuring DS-Lite Per Subnet Session Limitation to Prevent Denial of Service Attacks	243
Chapter 22	Transitioning to IPv6 Using 6rd Softwires	245
	Configuring a 6rd Softwire Concentrator	245
	Configuring Stateful Firewall Rules for 6rd Softwire	246
	Example: Basic 6rd Configuration	247
	Inter-Chassis High Availability for MS-MIC and MS-MPC	252
	Inter-Chassis High Availability for Stateful Firewall and NAPT44 Overview (MS-MIC, MS-MPC)	252
	Configuring Inter-Chassis High Availability for Stateful Firewall and NAPT44 (MS-MPC, MS-MIC)	253
	Example: Inter-Chassis Stateful High Availability for NAT and Stateful Firewall (MS-MIC, MS-MPC)	254
	High Availability and Load Balancing for 6rd Softwires	264
	Load Balancing a 6rd Domain Across Multiple Services PICs	264
	Example: Load Balancing a 6rd Domain Across Multiple Services PICs	264
	Configuring High Availability for 6rd Using 6rd Anycast	269
Chapter 23	Monitoring and Troubleshooting Softwires	271
	Ping and Traceroute for DS-Lite	271
	Monitoring Softwire Statistics	271
	Monitoring CGN, Stateful Firewall, and Softwire Flows	273
Part 5	Enabling Traffic to Pass Securely Using ALGs	
Chapter 24	ALG Overview	277
	ALG Descriptions	277
	Supported ALGs	277
	ALG Support Details	278
	Basic TCP ALG	279
	Basic UDP ALG	279
	BOOTP	280
	DCE RPC Services	280
	DNS	280
	FTP	280
	H323	281
	ICMP	282
	IIOP	282
	IP	282
	NetBIOS	282

	NetShow	283
	ONC RPC Services	283
	PPTP	283
	RealAudio	283
	Sun RPC and RPC Portmap Services	284
	RTSP	285
	SIP	286
	SNMP	286
	SQLNet	286
	TFTP	287
	Traceroute	287
	UNIX Remote-Shell Services	287
	Winframe	288
	Juniper Networks Defaults	288
	Examples: Referencing the Preset Statement from the Junos Default Group	298
	ALGs Available by Default for Junos OS Address Aware NAT	300
Chapter 25	ALGs Configuration Overview	303
	Configuring Application Sets	303
	Configuring Application Protocol Properties	303
	Configuring an Application Protocol	304
	Configuring the Network Protocol	306
	Configuring the ICMP Code and Type	307
	Configuring Source and Destination Ports	309
	Configuring the Inactivity Timeout Period	312
	Configuring SIP	312
	SIP ALG Interaction with Network Address Translation	313
	Junos OS SIP ALG Limitations	319
	Configuring an SNMP Command for Packet Matching	320
	Configuring an RPC Program Number	320
	Configuring the TTL Threshold	320
	Configuring a Universal Unique Identifier	320
	Examples: Configuring Application Protocols	321
	ICMP, Ping, and Traceroute ALGs for MS-MICs and MS-MPCs	322
	Monitoring Port Control Protocol Operations	323
Part 6	Securing Content Using Junos Network Secure and IDS	
Chapter 26	Junos Network Secure Overview	327
	Junos Network Secure Overview	327
	Stateful Firewall Support for Application Protocols	328
	Stateful Firewall Anomaly Checking	328
Chapter 27	Junos Network Secure Configuration Overview	331
	Configuring Stateful Firewall Rules	331
	Configuring Match Direction for Stateful Firewall Rules	332
	Configuring Match Conditions in Stateful Firewall Rules	332

	Configuring Actions in Stateful Firewall Rules	334
	Configuring IP Option Handling	334
	Configuring Stateful Firewall Rule Sets	335
	Examples: Configuring Stateful Firewall Rules	335
	Example: BOOTP and Broadcast Addresses	338
	Example: Configuring Layer 3 Services and the Services SDK on Two PICs	339
	Example: Virtual Routing and Forwarding (VRF) and Service Configuration . . .	351
Chapter 28	IDS Configuration Overview	355
	Configuring IDS Rules	355
	Configuring Match Direction for IDS Rules	356
	Configuring Match Conditions in IDS Rules	357
	Configuring Actions in IDS Rules	358
	Configuring IDS Rule Sets	363
	Examples: Configuring IDS Rules	364
Chapter 29	Monitoring Junos Network Secure	367
	Monitoring Stateful Firewall Conversations	367
	Monitoring CGN, Stateful Firewall, and Software Flows	367
	Monitoring Global Stateful Firewall Statistics	368
Part 7	Creating Secure Tunnels Using Junos VPN Site Secure	
Chapter 30	Junos VPN Site Secure Overview	371
	Understanding Junos VPN Site Secure	371
	IPsec	371
	Security Associations	372
	IKE	372
	Comparison of IPsec on ES PICs and Junos VPN Site Secure on Multiservices Line Cards	372
	Authentication Algorithms	374
	Encryption Algorithms	374
	IPsec Protocols	376
	Supported IPsec and IKE Standards	378
	IPsec Terms and Acronyms	379
Chapter 31	Junos VPN Site Secure Configuration Overview	383
	Minimum Security Association Configurations	383
	Minimum Manual SA Configuration	383
	Minimum Dynamic SA Configuration	384
	Configuring Security Associations	385
	Configuring Manual Security Associations	385
	Configuring the Direction for IPsec Processing	386
	Configuring the Protocol for a Manual IPsec SA	387
	Configuring the Security Parameter Index	387
	Configuring the Auxiliary Security Parameter Index	388
	Configuring Authentication for a Manual IPsec SA	388

Configuring Encryption for a Manual IPsec SA	389
Configuring Dynamic Security Associations	390
Clearing Security Associations	390
Example: Configuring Manual SAs	391
Configuring IKE Proposals	405
Configuring the Authentication Algorithm for an IKE Proposal	406
Configuring the Authentication Method for an IKE Proposal	406
Configuring the Diffie-Hellman Group for an IKE Proposal	407
Configuring the Encryption Algorithm for an IKE Proposal	408
Configuring the Lifetime for an IKE SA	408
Example: Configuring an IKE Proposal	409
Configuring IKE Policies	409
Configuring the IKE Phase	411
Configuring the Mode for an IKE Policy	411
Configuring the Proposals in an IKE Policy	411
Configuring the Preshared Key for an IKE Policy	411
Configuring the Local Certificate for an IKE Policy	412
Configuring a Certificate Revocation List	413
Configuring the Description for an IKE Policy	413
Configuring Local and Remote IDs for IKE Phase 1 Negotiation	413
Enabling Invalid SPI Recovery	414
Example: Configuring an IKE Policy	414
Configuring IPsec Proposals	415
Configuring the Authentication Algorithm for an IPsec Proposal	416
Configuring the Description for an IPsec Proposal	418
Configuring the Encryption Algorithm for an IPsec Proposal	418
Configuring the Lifetime for an IPsec SA	418
Configuring the Protocol for a Dynamic SA	419
Configuring IPsec Policies	420
Configuring the Description for an IPsec Policy	420
Configuring Perfect Forward Secrecy	421
Configuring the Proposals in an IPsec Policy	421
IPsec Policy for Dynamic Endpoints	421
Example: Configuring an IPsec Policy	422
Configuring IPsec Rules	422
Configuring Match Direction for IPsec Rules	424
Configuring Match Conditions in IPsec Rules	424
Configuring Actions in IPsec Rules	426
Enabling IPsec Packet Fragmentation	427
Configuring Destination Addresses for Dead Peer Detection	427
Configuring or Disabling IPsec Anti-Replay	428
Enabling System Log Messages	429
Specifying the MTU for IPsec Tunnels	429
Configuring IPsec Rule Sets	429
Service Sets for IPsec Tunnels	430
Configuring IPsec Service Sets	430
Configuring the Local Gateway Address for IPsec Service Sets	431
IKE Addresses in VRF Instances	432
Configuring IKE Access Profiles for IPsec Service Sets	432

	Configuring Certification Authorities for IPsec Service Sets	433
	Configuring or Disabling Antireplay Service	433
	Clearing the Don't-Fragment Bit	434
	Configuring Passive-Mode Tunneling	435
	Configuring the Tunnel MTU Value	436
	Tracing Junos VPN Site Secure Operations	436
	Disabling IPsec Tunnel Endpoint in Traceroute	437
	Tracing IPsec PKI Operations	438
	Multitask Example: Configuring IPsec Services	438
	Configuring the IKE Proposal	439
	Configuring the IKE Policy (and Referencing the IKE Proposal)	439
	Configuring the IPsec Proposal	440
	Configuring the IPsec Policy (and Referencing the IPsec Proposal)	441
	Configuring the IPsec Rule (and Referencing the IKE and IPsec Policies)	441
	Configuring IPsec Trace Options	442
	Configuring the Access Profile (and Referencing the IKE and IPsec Policies)	443
	Configuring the Service Set (and Referencing the IKE Profile and the IPsec Rule)	444
Chapter 32	Enhancing Security with Static IPsec over VRF	447
	Example: Configuring Statically Assigned IPsec Tunnels over a VRF Instance	447
Chapter 33	Dynamically Assigning Tunnels Using Junos VPN Site Secure	455
	Configuring Dynamic Endpoints for IPsec Tunnels	455
	Authentication Process	456
	Implicit Dynamic Rules	456
	Reverse Route Insertion	457
	Configuring an IKE Access Profile	457
	Referencing the IKE Access Profile in a Service Set	459
	Configuring the Interface Identifier	459
	Default IKE and IPsec Proposals	460
	Example: Configuring Dynamically Assigned Policy Based Tunnels	461
	Example: Configuring IKE Dynamic SAs	466
	Example: IKE Dynamic SA Configuration with Digital Certificates	483
Chapter 34	Enabling IPsec for the Services SDK	507
	Configuring Junos VPN Site Secure or IPsec VPN	507
Part 8	Alleviating Congestion and Controlling Service Using CoS	
Chapter 35	Class of Service Overview	511
	Class of Service Overview	511
Chapter 36	Class of Service Configuration Overview	513
	Restrictions and Cautions for CoS Configuration on Services Interfaces	513
	Configuring CoS Rules	514
	Configuring Match Direction for CoS Rules	515
	Configuring Match Conditions In CoS Rules	515

	Configuring Actions in CoS Rules	516
	Configuring Application Profiles for Use as CoS Rule Actions	517
	Configuring Reflexive and Reverse CoS Rule Actions	518
	Example: Configuring CoS Rules	518
	Configuring CoS Rule Sets	519
	Examples: Configuring CoS on Services Interfaces	519
Chapter 37	Configuring Class of Service on LSQ Interfaces	523
	Configuring CoS Scheduling Queues on Logical LSQ Interfaces	523
	Configuring Scheduler Buffer Size	524
	Configuring Scheduler Priority	525
	Configuring Scheduler Shaping Rate	525
	Configuring Drop Profiles	525
	Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces	527
	Configuring Link Services and CoS on Services PICs	529
	Oversubscribing Interface Bandwidth on LSQ Interfaces	532
	Examples: Oversubscribing an LSQ Interface	535
	Configuring Guaranteed Minimum Rate on LSQ Interfaces	537
	Example: Configuring Guaranteed Minimum Rate	539
Part 9	Configuring Interface Redundancy and Bundling on LSQ Interfaces	
Chapter 38	Overview	543
	Layer 2 Service Package Capabilities and Interfaces	543
Chapter 39	Configuring Interface Redundancy with SONET APS and Virtual Interfaces	545
	Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS	545
	Configuring the Association between LSQ and SONET Interfaces	546
	Configuring SONET APS Interoperability with Cisco Systems FRF.16	547
	Restrictions on APS Redundancy for LSQ Interfaces	547
	Configuring LSQ Interface Redundancy in a Single Router Using SONET APS	548
	Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces	548
	Configuring Redundant Paired LSQ Interfaces	549
	Restrictions on Redundant LSQ Interfaces	550
	Configuring Link State Replication for Redundant Link PICs	551
	Examples: Configuring Redundant LSQ Interfaces for Failure Recovery	553
Chapter 40	Enabling Bundling on LSQ Interfaces	559
	Inline MLPPP for WAN Interfaces Overview	559
	Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces	561
	Configuring Multiclass MLPPP on LSQ Interfaces	562
	Enabling Inline LSQ Services	563
	Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using MLPPP	565
	Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP	568
	Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.16	571
	Example: Configuring an LSQ Interface as an NxT1 Bundle Using FRF.16	574

	Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.15	576
	Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI	577
	Example: Configuring an LSQ Interface for a Fractional T1 Interface Using MLPPP and LFI	580
	Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12	582
	Examples: Configuring an LSQ Interface for a Fractional T1 Interface Using FRF.12	585
	Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP	589
	Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12	591
	Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP	593
Part 10	Enabling Load Balancing and High Availability Using Multiservices Interfaces	
Chapter 41	Enabling Load Balancing and High Availability Using Multiservices Interfaces	599
	Understanding Aggregated Multiservices Interfaces	599
	Aggregated Multiservices Interface	599
	IPv6 Traffic on AMS Interfaces Overview	603
	Member Failure Options and High Availability Settings	604
	Configuring Load Balancing on AMS Infrastructure	605
	Configuring AMS Infrastructure	606
	Configuring High Availability	607
	Load Balancing Network Address Translation Flows	607
	Example: Configuring an Aggregated Multiservices Interface (AMS)	608
	Example: Configuring Next-Hop Style Services on an Aggregated Multiservices Interface	613
	Example: Configuring Static Source Translation on AMS Infrastructure	616
Part 11	Handling VoIP, HTTP, and Layer 2 Traffic	
Chapter 42	Handling VoIP Traffic Using Voice Services	621
	Voice Services Overview	621
	Configuring Services Interfaces for Voice Services	622
	Configuring the Logical Interface Address for the MLPPP Bundle	622
	Configuring Compression of Voice Traffic	623
	Configuring Delay-Sensitive Packet Interleaving	624
	Example: Configuring Compression of Voice Traffic	624
	Configuring Encapsulation for Voice Services	625
	Configuring Network Interfaces for Voice Services	626
	Configuring Voice Services Bundles with MLPPP Encapsulation	626
	Configuring the Compression Interface with PPP Encapsulation	626
	Examples: Configuring Voice Services	627

Chapter 43	Handling HTTP Traffic Using HTTP Content Manager (HCM)	631
	HTTP Content Manager (HCM)	631
	Configuring the HTTP-Manager Package on the Router	631
	HTTP URL Tracking and Policy Control for Client Requests	634
	Guidelines for Configuring HTTP URL Monitoring for Client Requests	635
	Configuring HTTP URL Tracking and Policy Control	635
Chapter 44	Tunneling PPP Packets Across a Network Using Layer 2 Tunneling	637
	Layer 2 Tunneling Protocol Overview	637
	L2TP Services Configuration Overview	638
	L2TP Minimum Configuration	639
	Configuring L2TP Tunnel Groups	641
	Configuring Access Profiles for L2TP Tunnel Groups	642
	Configuring the Local Gateway Address and PIC	642
	Configuring Window Size for L2TP Tunnels	643
	Configuring Timers for L2TP Tunnels	643
	Hiding Attribute-Value Pairs for L2TP Tunnels	644
	Configuring System Logging of L2TP Tunnel Activity	644
	Configuring the Identifier for Logical Interfaces that Provide L2TP Services	646
	Example: Configuring Multilink PPP on a Shared Logical Interface	646
	AS PIC Redundancy for L2TP Services	648
	Examples: Configuring L2TP Services	648
	Tracing L2TP Operations	652
Part 12	Configuring Application Aware Services Interfaces	
Chapter 45	Configuring Stateless, Rule-Based Services Using Application-Aware Access Lists	657
	AACL Overview	657
	Best-Effort Application Identification of DPI-Serviced Flows	658
	Features that Support Application-Level Filtering	659
	Best-Effort Application Determination	659
	APPID, AACL, and L-PDF Processing in Preconvergence Scenarios	659
	Prior to a Final or Best-Effort Application Identification	659
	Upon Best-Effort Application Identification	660
	While Application Identification Is on a Best-Effort Basis	660
	If a Flow Ends Before an Application Identification Is Made	660
	If a Flow Ends While Application Identification on a Best-Effort Basis	660
	Configuring AACL Rules	661
	Configuring Match Direction for AACL Rules	662
	Configuring Match Conditions in AACL Rules	662
	Configuring Actions in AACL Rules	664
	Logging AACL Flows Based on Application	665
	Example: Configuring AACL Rules	666
	Configuring AACL Rule Sets	666
	Configuring Logging of AACL Flows	667

Chapter 46	Grouping Applications Together Using APPID	669
	APPID Overview	669
	Best-Effort Application Identification of DPI-Serviced Flows	671
	Features that Support Application-Level Filtering	672
	Best-Effort Application Determination	672
	APPID, ACL, and L-PDF Processing in Preconvergence Scenarios	672
	Prior to a Final or Best-Effort Application Identification	672
	Upon Best-Effort Application Identification	673
	While Application Identification Is on a Best-Effort Basis	673
	If a Flow Ends Before an Application Identification Is Made	673
	If a Flow Ends While Application Identification on a Best-Effort Basis	673
	Defining an Application Identification	674
	Configuring APPID Rules	676
	Using Stateful Firewall Rules to Identify Data Sessions	677
	Configuring Application Profiles	679
	Configuring Application Groups	680
	Application Identification for Nested Applications	681
	Disabling Application Identification for Nested Applications	682
	Configuring Global APPID Properties	683
	Configuring APPID Support for Heuristics	684
	Configuring APPID Support for Unidirectional Traffic	685
	Configuring Automatic Download of Application Package Updates	686
	Tracing APPID Operations	686
	Configuring the APPID Log Filename	687
	Configuring the Number and Size of APPID Log Files	687
	Configuring Access to the Log File	687
	Configuring a Regular Expression for Lines to Be Logged	688
	Configuring the Tracing Flags	688
	Examples: Configuring Application Identification Properties	688
Chapter 47	Detecting Suspicious and Anomalous Network Traffic Using IDP	691
	IDP Overview	691
	Best-Effort Application Identification of DPI-Serviced Flows	693
	Features that Support Application-Level Filtering	693
	Best-Effort Application Determination	694
	APPID, ACL, and L-PDF Processing in Preconvergence Scenarios	694
	Prior to a Final or Best-Effort Application Identification	694
	Upon Best-Effort Application Identification	695
	While Application Identification Is on a Best-Effort Basis	695
	If a Flow Ends Before an Application Identification Is Made	695
	If a Flow Ends While Application Identification on a Best-Effort Basis	695
Chapter 48	Collecting Statistics and Tracking Data Using L-PDF	697
	L-PDF Overview	697
	Best-Effort Application Identification of DPI-Serviced Flows	699
	Features that Support Application-Level Filtering	699
	Best-Effort Application Determination	700

	APPID, ACL, and L-PDF Processing in Preconvergence Scenarios	700
	Prior to a Final or Best-Effort Application Identification	700
	Upon Best-Effort Application Identification	701
	While Application Identification Is on a Best-Effort Basis	701
	If a Flow Ends Before an Application Identification Is Made	701
	If a Flow Ends While Application Identification on a Best-Effort Basis	701
	Configuring Statistics Profiles	702
	Configuring an L-PDF Statistics Profile	703
	Configuring an ACL Statistics Profile	704
	Applying L-PDF Profiles to Service Sets	705
	Tracing L-PDF Operations	707
Part 13	Configuring Link and Multilink Services Interfaces	
Chapter 49	Overview	711
	Link and Multilink Services Overview	711
	Multilink and Link Services PICs Overview	714
	Multilink Interfaces on Channelized MICs Overview	715
	Multilink and Link Services Logical Interface Configuration Overview	717
	Default Settings for Multilink and Link Services Logical Interfaces	717
Chapter 50	Achieving Greater Bandwidth, Load Balancing, and Redundancy with Multilink Bundles	719
	Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links	719
	Configuring the Number of Bundles on Link Services PICs	720
	Configuring the Links in a Multilink or Link Services Bundle	721
	Example: Configuring a Link Services Interface with Two Links	722
Chapter 51	Configuring the Physical and Logical Interfaces in a Multilink Configuration	725
	Configuring Link Services Physical Interfaces	725
	Default Settings for Link Services Interfaces	726
	Configuring Encapsulation for Link Services Physical Interfaces	726
	Configuring Acknowledgment Timers on Link Services Physical Interfaces	727
	Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16	727
	Configuring Keepalives on Link Services Physical Interfaces	728
	Configuring Encapsulation for Multilink and Link Services Logical Interfaces	729
	Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces	730
	Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces	732
	Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces	733
	Configuring MRRU on Multilink and Link Services Logical Interfaces	734

	Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces	735
	Example: Configuring Link Interfaces on Channelized MICs	735
Chapter 52	Bundling Multiple PPP Links on a Single Link Using MLPPP	747
	Example: Configuring a Multilink Interface with MLPPP	747
	Example: Configuring a Multilink Interface with MLPPP over ATM 2 Interfaces	748
	Example: Configuring an MLPPP Bundle	750
	Example: Configuring a Link Services Interface with MLPPP	753
	Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces	754
	Enabling MLPPP Link Fragmentation and Interleaving	769
	Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces	773
	Configuring LFI with DLCI Scheduling	774
	Example: Configuring LFI with DLCI Scheduling	775
Chapter 53	Bundling Multiple Frame Relay DLCIs into a Single Link Using MLFR	777
	Configuring DLCIs on Link Services Logical Interfaces	777
	Configuring Point-to-Point DLCIs for MLFR FRF.16 and MLPPP Bundles	777
	Configuring Multicast-Capable DLCIs for MLFR FRF.16 Bundles	778
	Example: Configuring a Multilink Interface with MLFR FRF.15	778
	Example: Configuring Multilink Frame Relay FRF.16	779
	Example: Configuring Multilink Frame Relay FRF.15	783
	Example: Configuring a Link Services Interface with MLFR FRF.15	786
	Example: Configuring a Link Services PIC with MLFR FRF.16	787
	Example: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces	788
Chapter 54	Configuring Additional Services on Link Services Interfaces	801
	Configuring CoS on Link Services Interfaces	801
	CoS for Link Services Interfaces on M Series and T Series Routers	801
	Example: Configuring CoS on Link Services Interfaces	803
	Example: Configuring Link and Voice Services Interfaces with a Combination of Bundle Types	806
Part 14	Flow Monitoring and Flow Collection Services	
Chapter 55	Monitoring Traffic Using Active Flow Monitoring	815
	Active Flow Monitoring Overview	815
	Configuring Flow Monitoring	818
	Configuring Flow-Monitoring Interfaces	818
	Configuring Flow-Monitoring Properties	820
	Directing Traffic to Flow-Monitoring Interfaces	820
	Exporting Flows	821
	Configuring Time Periods when Flow Monitoring is Active and Inactive	821

	Example: Configuring Flow Monitoring	822
	Example: Configuring Active Monitoring on Logical Systems	823
	Configuring Services Interface Redundancy with Flow Monitoring	826
	Flow Offloading	827
Chapter 56	Monitoring Traffic Using Passive Flow Monitoring	829
	Passive Flow Monitoring Overview	829
	Enabling Passive Flow Monitoring	830
	Passive Flow Monitoring for MPLS Encapsulated Packets	832
	Removing MPLS Labels from Incoming Packets	833
	Example: Enabling IPv4 Passive Flow Monitoring	834
	Example: Enabling IPv6 Passive Flow Monitoring	836
Chapter 57	Processing and Exporting Multiple Records Using Flow Collection	839
	Flow Collection Overview	839
	Configuring Flow Collection	840
	Configuring Destination FTP Servers for Flow Records	840
	Configuring a Packet Analyzer	841
	Configuring File Formats	841
	Configuring Interface Mappings	842
	Configuring Transfer Logs	842
	Configuring Retry Attempts	843
	Sending cflowd Records to Flow Collector Interfaces	843
	Configuring Flow Collection Mode and Interfaces on Services PICs	844
Part 15	Flow Capture Services	
Chapter 58	Dynamically Capturing Packet Flows Using Junos Capture Vision	847
	Understanding Junos Capture Vision	847
	Junos Capture Vision Architecture	847
	Liberal Sequence Windowing	848
	Intercepting IPv6 Flows	849
	Configuring Junos Capture Vision	849
	Configuring the Capture Group	849
	Configuring the Content Destination	850
	Configuring the Control Source	851
	Configuring the DFC PIC Interface	852
	Configuring the Firewall Filter	853
	Configuring System Logging	853
	Configuring Tracing Options for Junos Capture Vision Events	854
	Configuring Thresholds	854
	Limiting the Number of Duplicates of a Packet	855
	Example: Configuring Junos Capture Vision	855
Chapter 59	Detecting Threats and Intercepting Flows Using Junos Packet Vision	859
	Understanding Junos Packet Vision	859
	Junos Packet Vision Architecture	860
	Configuring Junos Packet Vision	861
	Configuring the Junos Packet Vision Interface	861
	Strengthening Junos Packet Vision Security	862

	Restrictions on Junos Packet Vision Services	863
	Configuring FlowTapLite	864
	Examples: Configuring Junos Packet Vision	865
Part 16	Sampling, Discard Accounting, and Port Mirroring Services	
Chapter 60	Sampling Data Using Traffic Sampling and Discard Accounting	871
	Configuring Traffic Sampling	871
	Configuring Firewall Filter for Traffic Sampling	871
	Configuring Traffic Sampling on a Logical Interface	873
	Disabling Traffic Sampling	874
	Sampling Once	874
	Preserving Prerewrite ToS Value for Egress Sampled or Mirrored Packets	875
	Configuring Traffic Sampling Output	876
	Traffic Sampling Output Format	877
	Tracing Traffic Sampling Operations	878
	Traffic Sampling Examples	878
	Example: Sampling a Single SONET/SDH Interface	878
	Example: Sampling All Traffic from a Single IP Address	879
	Example: Sampling All FTP Traffic	880
	Sampling Instance Configuration	881
	Configuring Discard Accounting	883
Chapter 61	Sampling Data Using Inline Sampling	885
	Understanding Inline Active Flow Monitoring	885
	Inline Active Flow Monitoring	885
	Inline Active Flow Monitoring Limitations and Restrictions	886
	IPFIX and Version 9 Templates	887
	Fields Included in the IPFIX IPv4 Template	888
	Fields Included in the IPFIX IPv6 Template	888
	Fields Included in the Version 9 IPv4 Template	889
	Configuring Inline Active flow Monitoring	890
	Configuring Inline Active Flow Monitoring on MX80 Routers	894
Chapter 62	Sampling Data Using Flow Aggregation	897
	Understanding Flow Aggregation	897
	Enabling Flow Aggregation	898
	Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd	898
	Configuring Flow Aggregation to Use Version 9 Flow Templates	902
	Configuring the Traffic to Be Sampled	903
	Configuring the Version 9 Template Properties	904
	Customizing Template ID, Observation Domain ID, and Source ID for Version 9 flow Templates	905
	Restrictions	906
	Fields Included in Each Template Type	906
	MPLS Sampling Behavior	908
	Verification	908
	Examples: Configuring Version 9 Flow Templates	909

	Configuring Flow Aggregation to Use IPFIX Flow Templates	912
	Configuring the IPFIX Template Properties	913
	Restrictions	914
	Customizing Template ID, Observation Domain ID, and Source ID for IPFIX flow Templates	914
	Fields Included in the IPv4 Template	915
	Fields Included in the IPv6 Template	916
	Verification	916
	Example: Configuring an IPFIX Flow Templates and Flow Sampling	917
	Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows	918
	Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows	921
	Directing Replicated Flows to Multiple Flow Servers	926
	Directing Replicated Routing Engine–Based Sampling Flows to Multiple Servers	926
	Directing Replicated Version 9 Flow Aggregates to Multiple Servers	927
	Logging cflowd Flows Before Export	928
Chapter 63	Sending Packets for Analysis Using Port Mirroring	931
	Understanding Port Mirroring	931
	Configuring Port Mirroring	931
	Configuring Tunnels	934
	Port Mirroring with Next-Hop Groups	936
	Configuring Inline Port Mirroring	937
	Filter-Based Forwarding with Multiple Monitoring Interfaces	938
	Restrictions	938
	Configuring Port Mirroring on Services Interfaces	939
	Examples: Configuring Port Mirroring	940
	Defining a Next-Hop Group for Port Mirroring	948
	Example: Multiple Port Mirroring with Next-Hop Groups Configuration	949
Part 17	Real-Time Performance Monitoring and Video Monitoring Services	
Chapter 64	Monitoring Traffic Using Real-Time Performance Monitoring	957
	Real-Time Performance Monitoring Services Overview	957
	Configuring RPM Probes	959
	Configuring RPM Receiver Servers	963
	Limiting the Number of Concurrent RPM Probes	964
	Configuring RPM Timestamping	964
	Configuring TWAMP	968
	Configuring TWAMP Interfaces	969
	Configuring TWAMP Servers	969
	Configuring BGP Neighbor Discovery Through RPM	971
	Examples: Configuring BGP Neighbor Discovery Through RPM	973
	Tracing RPM Operations	975
	Configuring the RPM Log File Name	975
	Configuring the Number and Size of RPM Log Files	975
	Configuring Access to the Log File	976

	Configuring a Regular Expression for Lines to Be Logged	976
	Configuring the Trace Operations	976
	Examples: Configuring Real-Time Performance Monitoring	977
	Enabling RPM for the Junos OS extension-provider package	981
Chapter 65	Testing the Performance of Network Devices Using RFC 2544-Based Benchmarking	983
	RFC2544-Based Benchmarking Tests Overview	983
	Layer 2 RFC2544-Based Benchmarking Tests Overview	986
	Supported RFC2544-Based Benchmarking Statements on MX104 Routers	988
	Configuring an RFC 2544-Based Benchmarking Test	989
	Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a IPv4 Network	989
	Configuring a Test Name for an RFC 2544-Based Benchmarking Test for an Ethernet Pseudowire	991
	Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a Layer 2 E-LAN Service in Bridge Domain	992
	Example: Configuring an RFC 2544-Based Benchmarking Test for Layer 3 IPv4 Services	993
	Example: Configuring an RFC 2544-Based Benchmarking Test for UNI Direction of Ethernet Pseudowires	1001
	Example: Configuring an RFC 2544-Based Benchmarking Test for NNI Direction of Ethernet Pseudowires	1008
	Example: Configuring RFC2544-Based Benchmarking Tests for Layer 2 E-LAN Services in Bridge Domains	1016
Chapter 66	Tracking Streaming Media Traffic Using Inline Video Monitoring	1043
	Inline Video Monitoring Overview	1043
	Configuring Inline Video Monitoring	1045
	Configuring Media Delivery Indexing Criteria	1045
	Configuring Interface Flow Criteria	1047
	Inline Video Monitoring Syslog Messages	1047
Part 18	Sampling, Discard Accounting, and Port Mirroring Services	
Chapter 67	Sampling Data Using Traffic Sampling and Discard Accounting	1051
	Configuring Traffic Sampling	1051
	Configuring Firewall Filter for Traffic Sampling	1051
	Configuring Traffic Sampling on a Logical Interface	1053
	Disabling Traffic Sampling	1054
	Sampling Once	1054
	Preserving Prerewrite ToS Value for Egress Sampled or Mirrored Packets	1055
	Configuring Traffic Sampling Output	1056
	Traffic Sampling Output Format	1057
	Tracing Traffic Sampling Operations	1058
	Traffic Sampling Examples	1058
	Example: Sampling a Single SONET/SDH Interface	1058
	Example: Sampling All Traffic from a Single IP Address	1059

	Example: Sampling All FTP Traffic	1060
	Sampling Instance Configuration	1061
	Configuring Discard Accounting	1063
Chapter 68	Sampling Data Using Inline Sampling	1065
	Understanding Inline Active Flow Monitoring	1065
	Inline Active Flow Monitoring	1065
	Inline Active Flow Monitoring Limitations and Restrictions	1066
	IPFIX and Version 9 Templates	1067
	Fields Included in the IPFIX IPv4 Template	1068
	Fields Included in the IPFIX IPv6 Template	1068
	Fields Included in the Version 9 IPv4 Template	1069
	Configuring Inline Active flow Monitoring	1070
	Configuring Inline Active Flow Monitoring on MX80 Routers	1074
Chapter 69	Sampling Data Using Flow Aggregation	1077
	Understanding Flow Aggregation	1077
	Enabling Flow Aggregation	1078
	Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd	1078
	Configuring Flow Aggregation to Use Version 9 Flow Templates	1082
	Configuring the Traffic to Be Sampled	1083
	Configuring the Version 9 Template Properties	1084
	Customizing Template ID, Observation Domain ID, and Source ID for Version 9 flow Templates	1085
	Restrictions	1086
	Fields Included in Each Template Type	1086
	MPLS Sampling Behavior	1088
	Verification	1088
	Examples: Configuring Version 9 Flow Templates	1089
	Configuring Flow Aggregation to Use IPFIX Flow Templates	1092
	Configuring the IPFIX Template Properties	1093
	Restrictions	1094
	Customizing Template ID, Observation Domain ID, and Source ID for IPFIX flow Templates	1094
	Fields Included in the IPv4 Template	1095
	Fields Included in the IPv6 Template	1096
	Verification	1096
	Example: Configuring an IPFIX Flow Templates and Flow Sampling	1097
	Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows	1098
	Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows	1101
	Directing Replicated Flows to Multiple Flow Servers	1106
	Directing Replicated Routing Engine–Based Sampling Flows to Multiple Servers	1106
	Directing Replicated Version 9 Flow Aggregates to Multiple Servers	1107
	Logging cflowd Flows Before Export	1108

Chapter 70	Sending Packets for Analysis Using Port Mirroring	1111
	Understanding Port Mirroring	1111
	Configuring Port Mirroring	1111
	Configuring Tunnels	1114
	Port Mirroring with Next-Hop Groups	1116
	Configuring Inline Port Mirroring	1117
	Filter-Based Forwarding with Multiple Monitoring Interfaces	1118
	Restrictions	1118
	Configuring Port Mirroring on Services Interfaces	1119
	Examples: Configuring Port Mirroring	1120
	Example: Multiple Port Mirroring with Next-Hop Groups Configuration	1128
 Part 19	 Real-Time Performance Monitoring and Video Monitoring Services	
 Chapter 71	 Monitoring Traffic Using Real-Time Performance Monitoring	 1135
	Real-Time Performance Monitoring Services Overview	1135
	Configuring RPM Probes	1137
	Configuring RPM Receiver Servers	1141
	Limiting the Number of Concurrent RPM Probes	1142
	Configuring RPM Timestamping	1142
	Configuring TWAMP	1146
	Configuring TWAMP Interfaces	1147
	Configuring TWAMP Servers	1147
	Configuring BGP Neighbor Discovery Through RPM	1149
	Examples: Configuring BGP Neighbor Discovery Through RPM	1151
	Tracing RPM Operations	1153
	Configuring the RPM Log File Name	1153
	Configuring the Number and Size of RPM Log Files	1153
	Configuring Access to the Log File	1154
	Configuring a Regular Expression for Lines to Be Logged	1154
	Configuring the Trace Operations	1154
	Examples: Configuring Real-Time Performance Monitoring	1155
	Enabling RPM for the Junos OS extension-provider package	1159
 Chapter 72	 Testing the Performance of Network Devices Using RFC 2544-Based Benchmarking	 1161
	RFC2544-Based Benchmarking Tests Overview	1161
	Configuring an RFC 2544-Based Benchmarking Test	1163
	Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a IPv4 Network	1164
	Configuring a Test Name for an RFC 2544-Based Benchmarking Test for an Ethernet Pseudowire	1165
	Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a Layer 2 E-LAN Service in Bridge Domain	1167
	Example: Configuring an RFC 2544-Based Benchmarking Test for Layer 3 IPv4 Services	1168
	Example: Configuring an RFC 2544-Based Benchmarking Test for UNI Direction of Ethernet Pseudowires	1175

	Example: Configuring an RFC 2544-Based Benchmarking Test for NNI Direction of Ethernet Pseudowires	1183
Chapter 73	Tracking Streaming Media Traffic Using Inline Video Monitoring	1191
	Inline Video Monitoring Overview	1191
	Configuring Inline Video Monitoring	1193
	Configuring Media Delivery Indexing Criteria	1193
	Configuring Interface Flow Criteria	1195
	Inline Video Monitoring Syslog Messages	1195
Part 20	Tunnel Services	
Chapter 74	Overview	1199
	Tunnel Services Overview	1199
	Configuring Tunnel Interfaces on MX Series Routers	1202
	Configuring Tunnel Interfaces on T4000 Routers	1203
Chapter 75	Encapsulating One Protocol Over Another Using GRE Interfaces	1205
	GRE Keepalive Time Overview	1205
	Configuring GRE Keepalive Time	1205
	Configuring Keepalive Time and Hold time for a GRE Tunnel Interface	1206
	Display GRE Keepalive Time Configuration	1206
	Display Keepalive Time Information on a GRE Tunnel Interface	1207
	Enabling Fragmentation on GRE Tunnels	1208
Chapter 76	Encapsulating One IP Packet Over Another Using IP-IP Interfaces	1211
	Configuring IPv6-over-IPv4 Tunnels	1211
	Example: Configuring an IPv6-over-IPv4 Tunnel	1211
Chapter 77	Filtering Unicast Packets Through Multicast Tunnel Interfaces	1213
	Configuring Unicast Tunnels	1213
	Configuring a Key Number on GRE Tunnels	1215
	Enabling Fragmentation on GRE Tunnels	1216
	Specifying an MTU Setting for the Tunnel	1216
	Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header	1217
	Configuring Packet Reassembly	1217
	Examples: Configuring Unicast Tunnels	1218
	Restricting Tunnels to Multicast Traffic	1219
Chapter 78	Connecting Logical Systems Using Logical Tunnel Interfaces	1221
	Configuring Logical Tunnel Interfaces	1221
	Connecting Logical Systems	1221
	Example: Configuring Logical Tunnels	1222
	Redundant Logical Tunnels Overview	1224
	Redundant Logical Tunnel Configuration	1224
	Redundant Logical Tunnel Failure Detection and Failover	1225
	Configuring Redundant Logical Tunnels	1226
	Example: Configuring Redundant Logical Tunnels	1227
Chapter 79	Understanding Default PIM Tunnel Configurations	1237
	Configuring PIM Tunnels	1237

Chapter 80	Facilitating VRF Table Lookup Using Virtual Loopback Tunnel Interfaces	1239
	Configuring Virtual Loopback Tunnels for VRF Table Lookup	1239
	Configuring Tunnel Interfaces for Routing Table Lookup	1241
	Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup	1241
	Example: Virtual Routing and Forwarding (VRF) and Service Configuration . . .	1242
Chapter 81	Enabling a VPN to Travel Through a Non-MPLS Network Using Dynamic Tunnels	1245
	Configuring Dynamic Tunnels	1245
Part 21	Encryption Services	
Chapter 82	Overview	1249
	Encryption Overview	1249
	Configuring an ES Tunnel Interface for a Layer 3 VPN	1249
Chapter 83	Sending Encrypted Traffic Through Tunnels	1251
	Configuring Encryption Interfaces	1251
	Specifying the Security Association Name for Encryption Interfaces	1252
	Configuring the MTU for Encryption Interfaces	1252
	Example: Configuring an Encryption Interface	1252
	Configuring Filters for Traffic Transiting the ES PIC	1253
	Traffic Overview	1253
	Configuring the Security Association	1254
	Configuring an Outbound Traffic Filter	1255
	Example: Configuring an Outbound Traffic Filter	1255
	Applying the Outbound Traffic Filter	1256
	Example: Applying the Outbound Traffic Filter	1256
	Configuring an Inbound Traffic Filter	1256
	Example: Configuring an Inbound Traffic Filter	1257
	Applying the Inbound Traffic Filter to the Encryption Interface	1257
	Example: Applying the Inbound Traffic Filter to the Encryption Interface	1257
Chapter 84	Configuring Redundancy in Case of Service Failure	1259
	Configuring ES PIC Redundancy	1259
	Example: Configuring ES PIC Redundancy	1259
	Configuring IPsec Tunnel Redundancy	1260
Part 22	Configuration Statements and Operational Commands	
Chapter 85	Configuration Statements	1265
	General Services Interfaces Configuration Statements	1265
	address (Interfaces)	1266
	applications (Services ALGs)	1267
	applications (Services CoS)	1267
	applications (Services IDS)	1268
	applications (Services NAT)	1268
	applications (Services Stateful Firewall)	1269

close-timeout	1269
cpu-load-threshold	1270
facility-override	1271
host (Interfaces)	1272
inactivity-timeout	1272
interfaces	1273
log-prefix (Interfaces)	1273
next-hop-service	1274
open-timeout	1275
port (System Log Messages)	1275
rule-set (Services Stateful Firewall)	1276
service-set (Interfaces)	1276
service-set (Services)	1277
services (CoS)	1279
services (IDS)	1279
services (IPsec VPN)	1280
services (Hierarchy)	1280
services (Interfaces)	1281
services (NAT)	1281
services (L2TP)	1282
services (L2TP System Logging)	1282
services (Stateful Firewall)	1283
services (System Logging)	1284
services-options	1285
syslog (Interfaces)	1286
tcp-tickles	1286
Adaptive Services Configuration Statements	1287
[edit services application-identification] Hierarchy Level	1295
IPsec Hierarchy Level	1297
adaptive-services-pics	1300
address (Interfaces)	1301
address (Services NAT Pool)	1301
address-allocation	1302
address-range	1302
aggregation	1303
allow-ip-options	1304
allow-multicast	1305
allow-overlapping-nat-pools	1305
anti-replay-window-size (Services IPsec VPN)	1306
anti-replay-window-size (Services Service Set)	1307
app-mapping-timeout	1308
application	1309
application-protocol	1310
application-profile	1312
application-set	1313
application-sets (Services CoS)	1313
application-sets (Services IDS)	1314
application-sets (Services NAT)	1314
application-sets (Services Stateful Firewall)	1315

applications (Services ALGs)	1315
applications (Services CoS)	1316
applications (Services IDS)	1316
applications (Services NAT)	1317
applications (Services Stateful Firewall)	1317
authentication (Services IPsec VPN)	1318
authentication-algorithm (Services IKE)	1319
authentication-algorithm (Services IPsec)	1320
authentication-method (Services IPsec VPN)	1321
auxiliary-spi (Services IPsec VPN)	1322
backup-remote-gateway	1322
bundle	1323
by-destination	1323
by-pair	1324
by-source	1325
bypass-traffic-on-exceeding-flow-limits	1325
bypass-traffic-on-pic-failure	1326
cgn-pic	1326
cisco-interoperability	1327
class	1328
clear-dont-fragment-bit (Interfaces GRE Tunnels)	1329
clear-dont-fragment-bit (Services IPsec VPN)	1329
clear-dont-fragment-bit (Services Service Set)	1330
clear-ike-sas-on-pic-restart	1330
clear-ipsec-sas-on-pic-restart	1331
compression	1331
compression-device (Interfaces)	1332
copy-dont-fragment-bit (Services IPsec VPN)	1332
copy-dont-fragment-bit (Services Set)	1333
data (FTP)	1333
dead-peer-detection (Services IPsec VPN)	1334
description (Services IPsec VPN)	1334
destination-address (Services CoS)	1335
destination-address (Services IDS)	1335
destination-address (Services IPsec VPN)	1336
destination-address (Services NAT)	1336
destination-address (Services Stateful Firewall)	1337
destination-address-range (Services IDS)	1337
destination-address-range (Services NAT)	1338
destination-address-range (Services Stateful Firewall)	1338
destination-pool	1339
destination-port	1340
destination-port range	1341
destination-prefix (Services IDS)	1341
destination-prefix (Services NAT)	1342
destination-prefix-ipv6	1342
destination-prefix-list (Services CoS)	1343
destination-prefix-list (Services IDS)	1343
destination-prefix-list (Services NAT)	1344

destination-prefix-list (Services Stateful Firewall)	1344
destined-port	1345
deterministic-port-block-allocation	1346
dh-group	1347
dial-options	1348
direction	1349
dns-alg-pool	1349
dns-alg-prefix	1350
drop-member-traffic (Aggregated Multiservices)	1350
ds-lite	1351
dscp	1352
dynamic	1352
ecmp-alb	1353
ei-mapping-timeout	1354
eif-flow-limit	1354
enable-rejoin (aggregated Multiservices)	1355
encapsulation	1356
encryption	1357
encryption-algorithm (Services IPsec VPN)	1358
establish-tunnels	1359
f-max-period	1359
facility-override (Service Sets)	1360
facility-override (System Log Reporting)	1361
family (Aggregated Multiservices)	1361
family (Interfaces)	1362
family (Voice Services)	1363
force-entry	1364
forwarding-class (Services CoS)	1364
forwarding-class (Services CoS Fragmentation Properties)	1365
fragment-threshold (Forwarding Class Maps)	1366
fragment-threshold (Interfaces LSQ)	1367
fragmentation-map	1367
fragmentation-maps	1368
from (Services CoS)	1369
from (Services IDS)	1370
from (Services IPsec VPN)	1371
from (Services HCM)	1371
from (Services NAT)	1372
from (Services Stateful Firewall)	1373
ftp (Services CoS)	1374
hello-interval	1374
hide-avps	1375
high-availability-options (aggregated Multiservices)	1376
host (L2TP)	1377
host (service-set)	1378
host (Services HCM)	1379
hot-standby	1379
icmp-code	1380
icmp-type	1380

ids-rules	1381
ignore-entry	1381
ike	1382
ike-access-profile	1383
inactivity-timeout	1383
initiate-dead-peer-detection	1384
input (Interfaces)	1384
interface-service	1385
interfaces (Aggregated Multiservices)	1386
interfaces (Voice Services)	1387
interval	1387
ipsec (Services IPsec VPN)	1388
ipsec-inside-interface	1388
ipsec-vpn-options	1389
ipsec-vpn-rules	1389
ipv6-multicast-interfaces	1390
l2tp-access-profile	1390
learn-sip-register	1391
lifetime-seconds (Services IPsec VPN)	1391
link-layer-overhead	1392
load-balance	1392
load-balancing-options (Aggregated Multiservices)	1393
local-certificate (Services IPsec VPN)	1394
local-gateway (IPSec)	1395
local-gateway (L2TP LNS)	1395
local-id	1396
log-prefix (L2TP)	1396
log-prefix (Services)	1397
logging (Services)	1397
logging (Services IDS)	1398
lsq-failure-options	1398
manual	1399
many-to-one (Aggregated Multiservices)	1400
mapping-refresh	1401
mapping-timeout	1402
match-direction (Services CoS)	1402
match-direction (Services IDS)	1403
match-direction (Services IPsec VPN)	1403
match-direction (Services NAT)	1404
match-direction (Services Stateful Firewall)	1404
max-drop-flows	1405
max-flows	1406
maximum-contexts	1407
maximum-send-window	1407
member-failure-options (Aggregated Multiservices)	1408
member-interface (Aggregated Multiservices)	1410
message-rate-limit	1411
mlfr-uni-nni-bundles-inline	1412
mode (Services IPsec VPN)	1413

mss	1413
multi-link-layer-2-inline	1414
multilink-class	1414
multilink-max-classes	1415
nat-options	1415
nat-rules	1416
next-hop-service	1417
no-anti-replay (Services IPsec VPN)	1418
no-anti-replay (Services Service Set)	1418
no-fragmentation	1419
no-ipsec-tunnel-in-traceroute	1419
no-per-unit-scheduler	1420
no-termination-request	1420
no-translation	1421
output	1421
overload-pool	1422
overload-prefix	1422
passive-mode-tunneling	1423
per-unit-scheduler	1424
perfect-forward-secrecy (Services IPsec VPN)	1425
pgcp-rules	1426
policy (Services IKE)	1427
policy (Services IPsec VPN)	1428
pool	1429
port (Services NAT)	1430
port (Services Voice)	1432
port (System Log Messages)	1432
port-forwarding	1433
port-forwarding-mappings	1433
ports-per-session	1434
post-service-filter	1434
ppp-access-profile	1435
pre-shared-key (Services IKE)	1435
preserve-interface	1436
primary (Adaptive Services Interfaces)	1436
primary (Link Services IQ PIC Interfaces)	1437
proposal (Services IKE)	1437
proposal (Services IPsec VPN)	1438
proposals	1438
protocol (Applications)	1439
protocol (IPSec)	1440
ptsp-rules	1440
queues	1441
receive-window	1441
redistribute-all-traffic (Aggregated Multiservices)	1442
redundancy-options (Adaptive Services Interfaces)	1443
redundancy-options (Link Services IQ PIC Interfaces)	1443
(reflexive reverse)	1444
rejoin-timeout (Aggregated Multiservices)	1445

remote-gateway	1445
remote-id	1446
request-url	1446
retransmit-interval (Services)	1447
rpc-program-number	1447
rtp	1448
rule (Services CoS)	1449
rule (Services IDS)	1450
rule (Services IPsec VPN)	1452
rule (Services NAT)	1454
rule (Services Stateful Firewall)	1455
rule (Softwire)	1456
rule-set (Services CoS)	1456
rule-set (Services IDS)	1457
rule-set (Services IPsec VPN)	1457
rule-set (Services NAT)	1458
rule-set (Services Stateful Firewall)	1458
rule-set (Softwire)	1459
secondary (Adaptive Services Interfaces)	1459
secondary (Link Services IQ PIC Interfaces)	1460
secure-nat-mapping	1460
secured-port-block-allocation	1461
server (pcp)	1463
service	1464
service-domain	1465
service-filter (Interfaces)	1465
service-interface (Adaptive Services Interfaces)	1466
service-interface (L2TP Processing)	1466
service-set (Interfaces)	1467
service-set (Services)	1468
service-set-options	1470
session-limit	1471
set-dont-fragment-bit (Services Set)	1472
set-dont-fragment-bit (Services IPsec VPN)	1472
sip-call-hold-timeout	1473
sip	1474
snmp-command	1474
softwire-concentrator	1475
softwire-options	1476
softwire-rules	1476
source-address (Service Sets)	1477
source-address (Services CoS)	1477
source-address (Services IDS)	1478
source-address (Services IPsec VPN)	1478
source-address (Services NAT)	1479
source-address (Services Stateful Firewall)	1479
source-address-range (Services IDS)	1480
source-address-range (Services NAT)	1480
source-address-range (Services Stateful Firewall)	1481

source-pool	1481
source-port	1482
source-prefix (Services IDS)	1482
source-prefix (Services NAT)	1483
source-prefix-ipv6	1483
source-prefix-list (Services CoS)	1484
source-prefix-list (Services IDS)	1484
source-prefix-list (Services NAT)	1485
source-prefix-list (Services Stateful Firewall)	1485
spi	1486
stateful-firewall-rules	1486
syslog (Services CoS)	1487
syslog (Services IDS)	1487
syslog (Services IPsec VPN)	1488
syslog (Services L2TP)	1488
syslog (Services NAT)	1489
syslog (Services Service Set)	1490
syslog (Services Stateful Firewall)	1491
syn-cookie	1492
tcp-mss	1493
term (Services CoS)	1494
term (Services IDS)	1495
term (Services IPsec VPN)	1497
term (Services HCM)	1498
term (Services NAT)	1499
term (Services Stateful Firewall)	1500
then (Services CoS)	1501
then (Services HCM)	1501
then (Services IDS)	1502
then (Services IPsec VPN)	1503
then (Services NAT)	1504
then (Services Stateful Firewall)	1505
threshold (Services IPsec)	1506
threshold (Services Logging and SYN-Cookie Defenses)	1506
traceoptions (Security PKI)	1507
traceoptions (Services IPsec VPN)	1509
traceoptions (Services L2TP)	1511
traceoptions (Services Logging)	1515
translated	1517
trigger-link-failure	1517
translated-port	1518
translation-type	1519
trusted-ca	1520
ttl-threshold	1521
tunnel-group	1522
tunnel-mtu (Services IPsec VPN)	1523
tunnel-mtu (Services Service Set)	1524
tunnel-timeout	1525
url	1525

url-list	1526
url-rule	1526
url-rule-set	1527
unit (Aggregated Multiservices)	1527
unit (Interfaces)	1528
unit (Voice Services)	1529
uuid	1530
v6rd	1531
version (IKE)	1532
video (Application Profile)	1532
voice (Application Profile)	1533
warm-standby	1533
Application Aware Services Configuration Statements	1533
[edit services aacl] Hierarchy List	1536
[edit services application-identification] Hierarchy Level	1537
aac1-fields	1539
aac1-statistics-profile	1540
address	1541
application (Defining)	1542
application (Including in Rule)	1543
application-aware-access-list-fields	1544
application-group	1545
application-group-any	1545
application-groups (Services AAC1)	1546
application-groups (Services Application Identification)	1546
application-system-cache-timeout	1547
application-unknown	1547
applications (Services AAC1)	1547
applications (Services Application Identification)	1548
automatic	1548
bypass-traffic-on-exceeding-flow-limits	1549
chain-order	1549
context	1550
destination (Services)	1550
destination-address	1551
destination-address-range	1551
destination-prefix-list	1552
direction	1552
disable (APPID Application)	1553
disable (APPID Application Group)	1553
disable (APPID Port Mapping)	1553
disable-global-timeout-override	1554
download	1554
enable-asymmetric-traffic-processing	1555
enable-heuristics	1555
file	1556
from	1557
idle-timeout	1558
ignore-errors	1558

index (Applications)	1559
index (Nested Applications)	1559
inactivity-non-tcp-timeout	1560
inactivity-tcp-timeout	1560
ip	1561
local-policy-decision-function	1562
log (aACL)	1563
match-direction	1563
max-checked-bytes	1564
maximum-transactions	1564
member	1565
min-checked-bytes	1565
nested-application	1566
nested-application-settings	1567
no-application-identification	1567
no-application-system-cache	1568
no-clear-application-system-cache	1568
no-nested-application	1569
no-protocol-method	1569
no-signature-based	1570
order (Services Application Identification)	1570
pattern	1571
policy-decision-statistics-profile	1572
port-mapping	1573
port-range	1573
profile	1574
protocol	1574
rule (AACL Rule Set)	1575
rule (Application Identification)	1576
rule (Including in Rule Set)	1577
rule-set (Services AACL)	1577
rule-set (Services Application Identification)	1578
service-set-options	1578
statistics (System Services)	1579
support-uni-directional-traffic	1579
service-set (Services)	1580
services (AACL)	1582
services (Application Identification)	1582
session-timeout (Application Identification)	1583
session-timeout (Interfaces)	1583
signature	1584
signature-method-all-ports	1584
source	1585
source-address (AACL)	1585
source-address-range	1586
source-prefix-list (Services AACL)	1586
source-prefix-list (Services IDS)	1587
term	1588
then	1589

traceoptions (Application Identification)	1591
traceoptions (Services Local Policy Decision Function)	1593
type	1594
type-of-service	1595
url	1595
Link and Multilink Services Configuration Statements	1595
acknowledge-retries	1597
acknowledge-timer	1597
action-red-differential-delay	1598
address (Interfaces)	1598
bundle	1599
destination (Interfaces)	1600
disable-mlppp-inner-ppp-pfc	1601
dlci	1601
drop-timeout	1602
encapsulation (Logical Interface)	1603
encapsulation (Physical Interface)	1604
family	1605
fragment-threshold	1606
hello-timer	1606
interfaces	1607
interleave-fragments	1607
lmi-type	1608
minimum-links	1608
mlfr-uni-nni-bundle-options	1609
mrru	1610
mtu	1611
multicast-dlci	1611
n391	1612
n392	1613
n393	1614
red-differential-delay	1614
short-sequence	1615
t391	1615
t392	1616
unit (Interfaces)	1617
yellow-differential-delay	1618
Monitoring, Sampling, and Collection Services Configuration Statements . . .	1618
[edit forwarding-options] Hierarchy Level	1624
[edit interfaces] Hierarchy Level	1627
[edit services dynamic-flow-control] Hierarchy Level	1628
[edit services flow-collector] Hierarchy Level	1629
[edit services flow-monitoring] Hierarchy Level	1630
[edit services flow-tap] Hierarchy Level	1631
[edit services rpm] Hierarchy Level	1631
accounting	1634
address (Interfaces)	1635
address (Services Dynamic Flow Capture)	1635
aggregate-export-interval	1636

aggregation	1637
allowed-destinations	1638
analyzer-address	1638
analyzer-id	1639
archive-sites	1639
authentication-mode	1640
autonomous-system-type	1641
bgp	1642
capture-group	1643
cflowd (Discard Accounting)	1644
client-list	1645
collector	1645
content-destination	1646
control-source	1647
core-dump	1648
data-fill	1649
data-format	1649
data-size	1650
destination (Interfaces)	1651
destination-interface	1652
destination-ipv4-address (RFC 2544 Benchmarking)	1653
destination-mac-address (RFC2544 Benchmarking)	1653
destination-port	1654
destination-udp-port (RFC 2544 Benchmarking)	1655
destinations	1655
direction (RFC2544 Benchmarking)	1656
disable (Forwarding Options)	1657
dscp-code-point	1658
duplicates-dropped-periodicity	1659
dynamic-flow-capture	1660
engine-id (Forwarding Options)	1661
engine-type	1662
export-format	1663
extension-service	1664
family (Monitoring)	1665
family (RFC2544 Benchmarking)	1666
family (Sampling)	1667
file (Sampling)	1668
file (Trace Options)	1669
file-specification (File Format)	1669
file-specification (Interface Mapping)	1670
filename	1670
filename-prefix	1671
files	1671
filter	1672
flow-active-timeout	1673
flow-collector	1674
flow-export-destination	1675
flow-export-rate	1675

flow-inactive-timeout	1676
flow-server	1677
flow-table-size	1678
flow-tap	1679
ftp (Flow Collector Files)	1680
ftp (Transfer Log Files)	1681
g-duplicates-dropped-periodicity	1682
g-max-duplicates	1683
hard-limit	1683
hard-limit-target	1684
hardware-timestamp	1684
history-size	1685
host-outbound	1685
udp-tcp-port-swap (RFC 2544 Benchmarking)	1686
in-service (RFC2544 Benchmarking)	1686
inactivity-timeout (Services RPM)	1687
inline-jflow	1687
input (Port Mirroring)	1688
input (Sampling)	1688
input-interface-index	1689
input-packet-rate-threshold	1689
instance (Sampling)	1690
interface (Accounting or Sampling)	1691
interface (Services Flow Tap)	1692
interface-map	1692
interfaces (Services Dynamic Flow Capture)	1693
interfaces (Video Monitoring)	1694
ip-swap (RFC 2544 Benchmarking)	1695
ipv4-flow-table-size	1695
ipv4-template	1696
ipv6-flow-table-size	1696
ipv6-template	1697
label-position	1697
local-dump	1698
logical-system	1698
match	1699
max-connection-duration	1699
max-duplicates	1700
max-packets-per-second	1701
maximum-age	1701
maximum-connections	1702
maximum-connections-per-client	1703
maximum-packet-length	1704
maximum-sessions	1705
maximum-sessions-per-connection	1706
minimum-priority	1706
mode (RFC 2544 Benchmarking)	1707
monitoring	1708
moving-average-size	1709

mpls-ipv4-template	1709
mpls-template	1710
multiservice-options	1710
name-format	1711
next-hop (Forwarding Options)	1712
next-hop-group (Forwarding Options)	1713
no-filter-check	1713
no-remote-trace (Trace Options)	1714
no-syslog	1714
notification-targets	1715
observation-domain-id	1716
one-way-hardware-timestamp	1717
option-refresh-rate	1718
options-template-id	1719
output (Accounting)	1720
output (Monitoring)	1721
output (Port Mirroring)	1722
output (Sampling)	1723
output-interface-index	1724
passive-monitor-mode	1724
password (Flow Collector File Servers)	1725
password (Transfer Log File Servers)	1725
peer-as-billing-template	1726
pic-memory-threshold	1726
pop-all-labels	1727
port (Flow Monitoring)	1728
port (RPM)	1728
port (TWAMP)	1729
pre-rewrite-tos	1729
probe	1730
probe-count	1731
probe-interval	1732
probe-limit	1732
probe-server	1733
probe-type	1734
rate (Forwarding Options)	1735
receive-options-packets	1735
receive-ttl-exceeded	1736
reflect-mode (RFC2544 Benchmarking)	1737
required-depth	1738
retry (Services Flow Collector)	1739
retry-delay	1739
rfc2544-benchmarking	1740
routing-instance	1741
routing-instances	1742
rpm (Interfaces)	1742
rpm (Services)	1743
run-length	1745
sample-once	1745

sampling (Forwarding Options)	1746
sampling (Interfaces)	1748
server	1749
server-inactivity-timeout	1749
service-port	1750
service-type (RFC2544 Benchmarking)	1750
services (RPM)	1751
shared-key	1751
size	1752
soft-limit	1753
soft-limit-clear	1753
source-address (Forwarding Options)	1754
source-address (Services)	1755
source-addresses	1755
source-id	1756
source-ipv4-address (RFC 2544 Benchmarking)	1756
source-mac-address (RFC2544 Benchmarking)	1757
source-udp-port (RFC 2544 Benchmarking)	1757
stamp	1758
syslog	1758
target (Services RPM)	1759
tcp	1759
templates	1760
test	1762
tests (RFC 2544 Benchmarking)	1763
test-interface (RFC 2544 Benchmarking)	1764
test-interval	1765
test-name (RFC 2544 Benchmarking)	1766
thresholds	1767
traceoptions (Forwarding Options)	1768
traceoptions (RPM)	1769
transfer	1770
transfer-log-archive	1771
traps	1772
ttl	1773
twamp	1774
twamp-server	1775
template (Forwarding Options)	1775
template-id	1776
template-refresh-rate	1777
trio-flow-offload	1777
udp	1778
unit	1779
username (Services)	1780
variant	1780
version	1781
version9 (Forwarding Options)	1781
video-monitoring	1782
world-readable	1783

Tunnel and Encryption Services Configuration Statements	1783
address (Interfaces)	1785
allow-fragmentation	1785
backup-destination	1786
backup-interface	1786
copy-tos-to-outer-ip-header	1787
destination (Interfaces)	1788
destination (Routing Instance)	1789
destination (Tunnel Remote End)	1789
destination-networks	1790
do-not-fragment	1791
dynamic-tunnels	1792
es-options	1793
family	1794
filter	1795
hold-time (OAM)	1795
interfaces	1796
ipsec-sa	1796
keepalive-time	1797
key	1798
multicast-only	1798
peer-unit	1799
reassemble-packets	1799
redundancy-group (Interfaces)	1800
redundancy-group (Logical Tunnels)	1801
routing-instance	1802
routing-instances	1802
routing-options	1803
source	1803
source	1804
source-address	1804
ttl	1805
tunnel	1806
tunnel	1807
unit (Interfaces)	1808
unit (Interfaces)	1809
Chapter 86	
Operational Commands	1811
Adaptive Services Operational Commands	1811
clear services cos statistics	1815
clear services crtp statistics	1816
clear services ids	1817
clear services ids destination-table	1818
clear services ids pair-table	1819
clear services ids source-table	1820
clear services inline nat pool	1821
clear services inline nat statistics	1822
clear services ipsec-vpn certificates	1823
clear services ipsec-vpn ike security-associations	1824

clear services ipsec-vpn ipsec security-associations	1825
clear services ipsec-vpn ipsec statistics	1826
clear services l2tp destination	1827
clear services l2tp destination statistics	1828
clear services l2tp multilink	1829
clear services l2tp session	1830
clear services l2tp session statistics	1832
clear services l2tp tunnel	1834
clear services l2tp tunnel statistics	1836
clear services nat flows	1838
clear services nat mappings	1839
clear services nat mappings app	1841
clear services nat mappings eim	1842
clear services nat mappings pcp	1844
clear security pki ca-certificate	1846
clear security pki certificate-request	1847
clear security pki crl	1848
clear security pki key-pair	1849
clear security pki local-certificate	1850
clear services service-sets statistics integrity-drops	1851
clear services service-sets statistics packet-drops	1852
clear services service-sets statistics syslog	1853
clear services stateful-firewall flows	1854
clear services stateful-firewall sip-call	1856
clear services stateful-firewall sip-register	1859
clear services stateful-firewall statistics	1862
request interface (revert switchover) (Adaptive Services)	1863
request security pki ca-certificate enroll	1864
request security pki ca-certificate load	1865
request security pki ca-certificate verify	1866
request security pki crl load	1867
request security pki generate-certificate-request	1868
request security pki generate-key-pair	1870
request security pki local-certificate enroll	1871
request security pki local-certificate generate-self-signed	1873
request security pki local-certificate load	1874
request security pki local-certificate verify	1875
request services ipsec-vpn ipsec switch tunnel	1876
show interfaces (Adaptive Services)	1877
show interfaces (Link Services IQ)	1885
show interfaces (Redundant Adaptive Services)	1909
show interfaces (Redundant Link Services IQ)	1911
show interfaces load-balancing	1925
show interfaces redundancy	1928
show security pki ca-certificate	1930
show security pki certificate-request	1934
show security pki crl	1936
show security pki local-certificate	1938
show services cos statistics	1941

show services crtp	1944
show services crtp flows	1946
show services ids	1948
show services inline nat pool	1956
show services inline nat statistics	1957
show services ipsec-vpn certificates	1958
show services ipsec-vpn ike security-associations	1961
show services ipsec-vpn ipsec security-associations	1965
show services ipsec-vpn ipsec statistics	1969
show services link-services cpu-usage	1972
show services l2tp multilink	1976
show services l2tp radius	1980
show services l2tp session	1984
show services l2tp summary	1992
show services l2tp tunnel	1997
show services l2tp user	2003
show services nat ipv6-multicast-interfaces	2007
show services nat mappings	2009
show services nat pool	2014
show services pcp statistics	2019
show services service-sets cpu-usage	2022
show services service-sets memory-usage	2024
show services service-sets statistics packet-drops	2026
show services service-sets statistics syslog	2028
show services service-sets statistics tcp-mss	2031
show services service-sets summary	2032
show services software	2034
show services software flows	2035
show services software statistics	2038
show services stateful-firewall conversations	2044
show services stateful-firewall flow-analysis	2048
show services stateful-firewall flows	2052
show services stateful-firewall sip-call	2058
show services stateful-firewall sip-register	2063
show services stateful-firewall statistics	2067
show services stateful-firewall statistics application-protocol sip	2076
show services stateful-firewall subscriber-analysis	2079
Application Aware Services Operational Commands	2081
clear services application-aware-access-list statistics	2083
clear services application-identification application-system-cache	2084
clear services application-identification counter	2085
clear services flows ip-action	2086
clear services local-policy-decision-function statistics	2087
request services application-identification application	2088
request services application-identification group	2090
show services application-aware-access-list flows	2091
show services application-identification application-system-cache	2094
show services application-identification counter	2096
show services application-identification group	2099

show services application-aware-access-list statistics	2101
show services application-identification application	2103
show services application-identification version	2106
show services flows	2107
show services local-policy-decision-function flows	2114
show services local-policy-decision-function statistics	2116
Link and Multilink Services Operational Commands	2117
show interfaces (Link Services)	2118
show interfaces (Link Services IQ)	2131
show interfaces (Multilink Services)	2155
Monitoring, Sampling, and Collection Services Operational Commands	2162
clear passive-monitoring statistics	2164
clear services accounting statistics inline-jflow	2165
clear services dynamic-flow-capture	2166
clear services flow-collector statistics	2167
clear services rpm twamp server connection	2168
clear services video-monitoring mdi errors fpc-slot	2169
clear services video-monitoring mdi statistics fpc-slot	2170
request services flow-collector change-destination primary interface	2171
request services flow-collector change-destination secondary interface	2172
request services flow-collector test-file-transfer	2173
show forwarding-options next-hop-group	2174
show forwarding-options port-mirroring	2177
show interfaces (Dynamic Flow Capture)	2179
show interfaces (Flow Collector)	2183
show interfaces (Flow Monitoring)	2189
show passive-monitoring error	2194
show passive-monitoring flow	2196
show passive-monitoring memory	2198
show passive-monitoring status	2200
show passive-monitoring usage	2202
show services accounting aggregation	2204
show services accounting aggregation template	2208
show services accounting errors	2209
show services accounting flow	2213
show services accounting flow-detail	2218
show services accounting memory	2223
show services accounting packet-size-distribution	2225
show services accounting status	2227
show services accounting usage	2230
show services dynamic-flow-capture content-destination	2232
show services dynamic-flow-capture control-source	2234
show services dynamic-flow-capture statistics	2236
show services flow-collector file interface	2239
show services flow-collector input interface	2241
show services flow-collector interface	2243
show services rpm active-servers	2249
show services rpm history-results	2250
show services rpm probe-results	2253

show services rpm rfc2544-benchmarking	2259
show services rpm rfc2544-benchmarking test-id	2264
show services rpm twamp server connection	2281
show services rpm twamp server session	2283
show services video-monitoring mdi errors fpc-slot	2285
show services video-monitoring mdi flows fpc-slot	2287
show services video-monitoring mdi stats fpc-slot	2291
test services rpm rfc2544-benchmarking test	2293
Tunnel and Encryption Services Operational Commands	2294
clear ike security-associations	2296
clear ipsec security-associations	2297
request ipsec switch	2299
request security certificate (signed)	2300
request security certificate (unsigned)	2302
request security key-pair	2303
request system certificate add	2304
show ike security-associations	2305
show interfaces (Encryption)	2309
show interfaces (GRE)	2315
show interfaces (IP-over-IP)	2322
show interfaces (Logical Tunnel)	2326
show interfaces (Multicast Tunnel)	2331
show interfaces (PIM)	2336
show interfaces (Virtual Loopback Tunnel)	2340
show ipsec certificates	2345
show ipsec redundancy	2348
show ipsec security-associations	2350
show system certificate	2353

Part 23

Index

Index	2357
-----------------	------

List of Figures

Part 2	Adaptive Services Overview	
Chapter 3	Adaptive Services Overview	25
	Figure 1: Packet Flow Through the Adaptive Services or MultiServices PIC	28
Part 3	Translating IP Addresses Using NAT	
Chapter 5	NAT Overview	55
	Figure 2: Dynamic NAT Flow	59
	Figure 3: Stateful NAT64 Flow	59
	Figure 4: DS-Lite Flow	60
Chapter 7	Avoiding IPv4 Exhaustion Using Junos Address Aware Network Addressing and Stateful NAT64	87
	Figure 5: IPv4 Depletion Solution - IPv4 Access Network	88
	Figure 6: IPv4 Depletion Solution - IPv6 Access Network	88
Chapter 8	Hiding Private Networks Using Static Source NAT	93
	Figure 7: Configuring NAT for Multicast Traffic	105
Chapter 11	Securing Traffic Using NAT-PT and ALGs	141
	Figure 8: Configuring DNS ALGs with NAT-PT Network Topology	151
Chapter 12	Reducing Traffic and Bandwidth Requirements Using Port Control Protocol	165
	Figure 9: Basic PCP NAPT44 Topology	166
	Figure 10: PCP with DS-Lite Plain Mode	166
	Figure 11: PCP with DS-Lite Tunnel Mode	167
	Figure 12: PCP with NAPT44	170
Chapter 16	Achieving Line-Rate, Low-Latency Translations Using Inline NAT	199
	Figure 13: Supported Inline NAT Types	200
	Figure 14: Deploy Inline NAT within L3VPN	202
Part 4	Transitioning to IPv6 Using Softwires	
Chapter 18	Softwires Overview	213
	Figure 15: 6rd Softwire Flow	216
Chapter 21	Transitioning to IPv6 Using DS-Lite Softwires	227
	Figure 16: DS-Lite Topology	229
Chapter 22	Transitioning to IPv6 Using 6rd Softwires	245
	Figure 17: Inter-Chassis High Availability Topology	253

	Figure 18: Inter-Chassis High Availability Topology	255
Part 7	Creating Secure Tunnels Using Junos VPN Site Secure	
Chapter 30	Junos VPN Site Secure Overview	371
	Figure 19: AH Protocol	376
	Figure 20: ESP Protocol	377
Chapter 31	Junos VPN Site Secure Configuration Overview	383
	Figure 21: Manual SA Topology	392
Chapter 33	Dynamically Assigning Tunnels Using Junos VPN Site Secure	455
	Figure 22: IPsec Dynamic Endpoint Tunneling Topology	462
	Figure 23: IKE Dynamic SAs	467
	Figure 24: MS PIC IKE Dynamic SA Topology Diagram	484
Part 9	Configuring Interface Redundancy and Bundling on LSQ Interfaces	
Chapter 40	Enabling Bundling on LSQ Interfaces	559
	Figure 25: Inline MLPPP for WAN Interfaces	560
Part 13	Configuring Link and Multilink Services Interfaces	
Chapter 50	Achieving Greater Bandwidth, Load Balancing, and Redundancy with Multilink Bundles	719
	Figure 26: Multilink Interface Configuration	721
Chapter 52	Bundling Multiple PPP Links on a Single Link Using MLPPP	747
	Figure 27: Configuring MLPPP and LFI on Serial Links	750
	Figure 28: Configuring Inline MLPPP and Multilink Frame Relay End-End (FRF.15) for WAN Interfaces	755
Chapter 53	Bundling Multiple Frame Relay DLCIs into a Single Link Using MLFR	777
	Figure 29: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces	789
Part 14	Flow Monitoring and Flow Collection Services	
Chapter 55	Monitoring Traffic Using Active Flow Monitoring	815
	Figure 30: Active Monitoring Configuration Topology	817
Chapter 56	Monitoring Traffic Using Passive Flow Monitoring	829
	Figure 31: Passive Monitoring Application Topology	830
Part 15	Flow Capture Services	
Chapter 58	Dynamically Capturing Packet Flows Using Junos Capture Vision	847
	Figure 32: Junos Capture Vision Topology	848
Chapter 59	Detecting Threats and Intercepting Flows Using Junos Packet Vision	859
	Figure 33: Junos Packet Vision Topology	861

Part 16	Sampling, Discard Accounting, and Port Mirroring Services	
Chapter 60	Sampling Data Using Traffic Sampling and Discard Accounting	871
	Figure 34: Configuring Sampling Rate	873
Chapter 63	Sending Packets for Analysis Using Port Mirroring	931
	Figure 35: Active Flow Monitoring—Multiple Port Mirroring with Next-Hop Groups Topology Diagram	950
Part 17	Real-Time Performance Monitoring and Video Monitoring Services	
Chapter 65	Testing the Performance of Network Devices Using RFC 2544-Based Benchmarking	983
	Figure 36: E-LAN and E-Line Reflection in Metro Solution	986
	Figure 37: RFC 2544-Based Benchmarking Test for a Layer 3 IPv4 Service	994
	Figure 38: RFC 2544-Based Benchmarking Test for UNI Direction of an Ethernet Pseudowire	1002
	Figure 39: RFC 2544-Based Benchmarking Test for NNI Direction of an Ethernet Pseudowire	1010
	Figure 40: Layer 2 reflection Simple Topology	1017
Part 18	Sampling, Discard Accounting, and Port Mirroring Services	
Chapter 67	Sampling Data Using Traffic Sampling and Discard Accounting	1051
	Figure 41: Configuring Sampling Rate	1053
Chapter 70	Sending Packets for Analysis Using Port Mirroring	1111
	Figure 42: Active Flow Monitoring—Multiple Port Mirroring with Next-Hop Groups Topology Diagram	1129
Part 19	Real-Time Performance Monitoring and Video Monitoring Services	
Chapter 72	Testing the Performance of Network Devices Using RFC 2544-Based Benchmarking	1161
	Figure 43: RFC 2544-Based Benchmarking Test for a Layer 3 IPv4 Service	1169
	Figure 44: RFC 2544-Based Benchmarking Test for UNI Direction of an Ethernet Pseudowire	1176
	Figure 45: RFC 2544-Based Benchmarking Test for NNI Direction of an Ethernet Pseudowire	1184
Part 20	Tunnel Services	
Chapter 78	Connecting Logical Systems Using Logical Tunnel Interfaces	1221
	Figure 46: Redundant Logical Tunnels	1224
	Figure 47: Redundant Logical Tunnels	1228
Part 21	Encryption Services	
Chapter 83	Sending Encrypted Traffic Through Tunnels	1251

	Figure 48: Example: IPsec Tunnel Connecting Security Gateways	1253
Chapter 84	Configuring Redundancy in Case of Service Failure	1259
	Figure 49: IPsec Tunnel Redundancy	1260

List of Tables

	About the Documentation	lvii
	Table 1: Notice Icons	lix
	Table 2: Text and Syntax Conventions	lx
Part 1	Services Interfaces Overview	
Chapter 1	Overview	3
	Table 3: MX Series Routers that Support MS-MIC and MS-MPC	6
Chapter 2	Configuration Overview	9
	Table 4: AS and Multiservices PIC Services by Service Package, PIC, and Platform	12
	Table 5: System Log Message Severity Levels	21
Part 2	Adaptive Services Overview	
Chapter 4	Adaptive Services Configuration Overview	29
	Table 6: System Log Message Severity Levels	47
	Table 7: Adaptive Services Tracing Flags	51
Part 3	Translating IP Addresses Using NAT	
Chapter 5	NAT Overview	55
	Table 8: Carrier-Grade NAT—Feature Comparison by Platform	61
	Table 9: Carrier-Grade NAT Translation Types	63
Chapter 10	Allowing Components of a Private Network to Share a Single Address Using NAPT	117
	Table 10: Deterministic Port Block Allocation Commit Constraints	126
	Table 11: Comparison of NAPT Implementation Methods	126
Chapter 11	Securing Traffic Using NAT-PT and ALGs	141
	Table 12: ALGs Available by Default	141
Part 5	Enabling Traffic to Pass Securely Using ALGs	
Chapter 24	ALG Overview	277
	Table 13: ALGs Supported by Junos OS	277
	Table 14: RealAudio Product Port Usage	283
	Table 15: Supported RPC Services	284
	Table 16: ALGs Available by Default	300
Chapter 25	ALGs Configuration Overview	303

	Table 17: Application Protocols Supported by Services Interfaces	304
	Table 18: Network Protocols Supported by Services Interfaces	306
	Table 19: ICMP Codes and Types Supported by Services Interfaces	308
	Table 20: Port Names Supported by Services Interfaces	309
	Table 21: Requesting Messages with NAT Table	317
Part 6	Securing Content Using Junos Network Secure and IDS	
Chapter 27	Junos Network Secure Configuration Overview	331
	Table 22: IP Option Values	335
Part 7	Creating Secure Tunnels Using Junos VPN Site Secure	
Chapter 30	Junos VPN Site Secure Overview	371
	Table 23: Statement Equivalents for ES and AS Interfaces	373
Chapter 33	Dynamically Assigning Tunnels Using Junos VPN Site Secure	455
	Table 24: Default IKE and IPsec Proposals for Dynamic Negotiations	460
Part 10	Enabling Load Balancing and High Availability Using Multiservices Interfaces	
Chapter 41	Enabling Load Balancing and High Availability Using Multiservices Interfaces	599
	Table 25: Key Configuration Statements Used in this Example	611
Part 11	Handling VoIP, HTTP, and Layer 2 Traffic	
Chapter 44	Tunneling PPP Packets Across a Network Using Layer 2 Tunneling	637
	Table 26: System Log Message Severity Levels	644
Part 13	Configuring Link and Multilink Services Interfaces	
Chapter 49	Overview	711
	Table 27: Multilink and Link Services PIC Capacities	714
	Table 28: Multilink and Link Services Logical Interface Statements	717
Chapter 50	Achieving Greater Bandwidth, Load Balancing, and Redundancy with Multilink Bundles	719
	Table 29: Link Services Bundle	722
Chapter 51	Configuring the Physical and Logical Interfaces in a Multilink Configuration	725
	Table 30: Link Services Physical Interface Statements for MLFR FRF.16	726
Chapter 54	Configuring Additional Services on Link Services Interfaces	801
	Table 31: Link Services CoS Queues	801
Part 16	Sampling, Discard Accounting, and Port Mirroring Services	
Chapter 62	Sampling Data Using Flow Aggregation	897
	Table 32: Example of Observation Domain ID	919

	Table 33: Values of Template and Option Template IDs for IPFIX Flows	923
	Table 34: Values of Template and Option Template IDs for Version 9 Flows . . .	923
	Table 35: Values of Template and Option Template IDs for IPFIX Flows	924
Part 17	Real-Time Performance Monitoring and Video Monitoring Services	
Chapter 64	Monitoring Traffic Using Real-Time Performance Monitoring	957
	Table 36: RPM Tracing Flags	976
Chapter 65	Testing the Performance of Network Devices Using RFC 2544-Based Benchmarking	983
	Table 37: Supported Network Topologies for RFC2544 Benchmarking Tests . .	984
	Table 38: Supported Interfaces for RFC2544 Benchmarking Tests	985
	Table 39: MAC Address Swapping Behavior for E-LAN and E-Line Services . . .	987
	Table 40: Supported RFC2544-Based Benchmarking Reflector Statements on MX104	988
Chapter 66	Tracking Streaming Media Traffic Using Inline Video Monitoring	1043
	Table 41: MPC Flow Monitoring Capacity by Model	1045
Part 18	Sampling, Discard Accounting, and Port Mirroring Services	
Chapter 69	Sampling Data Using Flow Aggregation	1077
	Table 42: Example of Observation Domain ID	1099
	Table 43: Values of Template and Option Template IDs for IPFIX Flows	1103
	Table 44: Values of Template and Option Template IDs for Version 9 Flows . . .	1103
	Table 45: Values of Template and Option Template IDs for IPFIX Flows	1104
Part 19	Real-Time Performance Monitoring and Video Monitoring Services	
Chapter 71	Monitoring Traffic Using Real-Time Performance Monitoring	1135
	Table 46: RPM Tracing Flags	1154
Chapter 72	Testing the Performance of Network Devices Using RFC 2544-Based Benchmarking	1161
	Table 47: Supported Network Topologies for RFC2544 Benchmarking Tests . .	1162
	Table 48: Supported Interfaces for RFC2544 Benchmarking Tests	1163
Chapter 73	Tracking Streaming Media Traffic Using Inline Video Monitoring	1191
	Table 49: MPC Flow Monitoring Capacity by Model	1193
Part 20	Tunnel Services	
Chapter 74	Overview	1199
	Table 50: Tunnel Interface Types	1199
Chapter 80	Facilitating VRF Table Lookup Using Virtual Loopback Tunnel Interfaces	1239
	Table 51: Methods for Configuring Egress Filtering	1239

Part 22	Configuration Statements and Operational Commands	
Chapter 85	Configuration Statements	1265
	Table 52: Behavior of Member Interface After One Multiservices PIC Fails	1408
	Table 53: Behavior of Member Interface After Two Multiservices PICs Fail	1409
Chapter 86	Operational Commands	1811
	Table 54: clear services nat flows Output Fields	1838
	Table 55: clear services nat mappings Output Fields	1839
	Table 56: clear services nat mappings app Output Fields	1841
	Table 57: clear services nat mappings eim Output Fields	1842
	Table 58: clear services nat mappings pcp Output Fields	1844
	Table 59: clear services stateful-firewall flows Output Fields	1855
	Table 60: clear services stateful-firewall sip-call Output Fields	1858
	Table 61: clear services stateful-firewall sip-register Output Fields	1861
	Table 62: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields	1877
	Table 63: show interfaces (Link Services IQ) Output Fields	1886
	Table 64: show interfaces (Redundant Link Services IQ) Output Fields	1911
	Table 65: Aggregated Multiservices show interfaces load-balancing Output Fields	1925
	Table 66: show interfaces redundancy Output Fields	1928
	Table 67: show security pki ca-certificate Output Fields	1930
	Table 68: show security pki certificate-request Output Fields	1934
	Table 69: show security pki crl Output Fields	1936
	Table 70: show security pki local-certificate Output Fields	1938
	Table 71: show services cos statistics Output Fields	1941
	Table 72: show services crtp Output Fields	1944
	Table 73: show services crtp flows Output Fields	1946
	Table 74: show services ids Output Fields	1949
	Table 75: show services inline nat pool Output Fields	1956
	Table 76: show services inline nat statistics Output Fields	1957
	Table 77: show services ipsec-vpn certificates Output Fields	1958
	Table 78: show services ipsec-vpn ike security-associations Output Fields	1961
	Table 79: show services ipsec-vpn ipsec security-associations Output Fields . .	1965
	Table 80: show services ipsec-vpn ipsec statistics Output Fields	1969
	Table 81: show services link-services cpu-usage Output Fields	1972
	Table 82: show services l2tp multilink Output Fields	1976
	Table 83: show services l2tp radius Output Fields	1980
	Table 84: show services l2tp session Output Fields	1985
	Table 85: show services l2tp summary Output Fields	1992
	Table 86: show services l2tp tunnel Output Fields	1998
	Table 87: show services l2tp user Output Fields	2003
	Table 88: show services nat ipv6-multicast-interfaces Output Fields	2007
	Table 89: show services nat mappings Output Fields	2010
	Table 90: show services nat pool Output Fields	2015
	Table 91: show services pcp statistics Output Fields	2019
	Table 92: show services service-sets cpu-usage Output Fields	2022
	Table 93: show services service-sets memory-usage Output Fields	2024
	Table 94: show services service-sets packet-drops Output Fields	2026

Table 95: show services service-sets statistics syslog Output Fields	2028
Table 96: show services service-sets statistics tcp-mss Output Fields	2031
Table 97: show services service-sets summary Output Fields	2032
Table 98: show-services-softwire Output Fields	2034
Table 99: show services softwire flows Output Fields	2036
Table 100: command-name Output Fields	2038
Table 101: show services stateful-firewall conversations Output Fields	2046
Table 102: show services stateful-firewall flow-analysis Output Fields	2048
Table 103: show services stateful-firewall flows Output Fields	2054
Table 104: show services stateful-firewall sip-call Output Fields	2060
Table 105: show services stateful-firewall sip-register Output Fields	2065
Table 106: show services stateful-firewall statistics Output Fields	2067
Table 107: show services stateful-firewall statistics application-protocol-sip Output Fields	2076
Table 108: show services stateful-firewall subscriber-analysis Output Fields . .	2079
Table 109: show services application-aware-access-list flows Output Fields . .	2091
Table 110: show application-identification application-system-cache Output Fields	2094
Table 111: show services application-identification counter Output Fields . . .	2096
Table 112: show services application-identification group Output Fields	2099
Table 113: show services application-aware-access-list statistics Output Fields	2101
Table 114: show services application-identification application Output Fields . .	2103
Table 115: show services flows Output Fields	2109
Table 116: show services local-policy-decision-function flows Output Fields . .	2114
Table 117: show services local-policy-decision-function statistics Output Fields	2116
Table 118: Link Services show interfaces Output Fields	2118
Table 119: show interfaces (Link Services IQ) Output Fields	2132
Table 120: Multilink Services show interfaces Output Fields	2155
Table 121: show forwarding-options next-hop-group Output Fields	2174
Table 122: show forwarding-options port-mirroring Output Fields	2177
Table 123: Dynamic Flow Capture show interfaces Output Fields	2179
Table 124: Flow Collector Show interfaces Output Fields	2183
Table 125: show interfaces Output Fields (Flow Monitoring)	2189
Table 126: show passive-monitoring error Output Fields	2194
Table 127: show passive-monitoring flow Output Fields	2196
Table 128: show passive-monitoring memory Output Fields	2198
Table 129: show passive-monitoring status Output Fields	2200
Table 130: show passive-monitoring usage Output Fields	2202
Table 131: show services accounting aggregation Output Fields	2205
Table 132: show services accounting aggregation template Output Fields . . .	2208
Table 133: show services accounting errors Output Fields	2209
Table 134: show services accounting flow Output Fields	2213
Table 135: show services accounting flow-detail Output Fields	2219
Table 136: show services accounting memory Output Fields	2223
Table 137: show services accounting packet-size-distribution Output Fields . .	2225
Table 138: show services accounting status Output Fields	2227
Table 139: show services accounting usage Output Fields	2230

Table 140: show services dynamic-flow-capture content-destination Output Fields	2232
Table 141: show services dynamic-flow-capture control-source Output Fields	2234
Table 142: show services dynamic-flow-capture statistics Output Fields	2236
Table 143: show services flow-collector file interface Output Fields	2239
Table 144: show services flow-collector input interface Output Fields	2241
Table 145: show services flow-collector interface Output Fields	2243
Table 146: show services rpm active-servers Output Fields	2249
Table 147: show services rpm history-results Output Fields	2250
Table 148: show services rpm probe-results Output Fields	2253
Table 149: show services rpm rfc2544-benchmarking Output Fields	2260
Table 150: show services rpm rfc2544-benchmarking test-id Output Fields	2265
Table 151: show services rpm twamp server connection Output Fields	2281
Table 152: show services rpm twamp server session Output Fields	2283
Table 153: show services video-monitoring mdi errors fpc-slot Output Fields	2285
Table 154: show services mdi flows Output Fields	2288
Table 155: show services video-monitoring mdi stats fpc-slot Output Fields	2291
Table 156: show ike security-associations Output Fields	2305
Table 157: Encryption show interfaces Output Fields	2309
Table 158: GRE show interfaces Output Fields	2316
Table 159: IP-over-IP show interfaces Output Fields	2322
Table 160: Logical Tunnel show interfaces Output Fields	2326
Table 161: Multicast Tunnel show interfaces Output Fields	2332
Table 162: PIM show interfaces Output Fields	2336
Table 163: Virtual Loopback Tunnel show interfaces Output Fields	2340
Table 164: show ipsec certificates Output Fields	2345
Table 165: show ipsec redundancy Output Fields	2348
Table 166: show ipsec security-associations Output Fields	2350
Table 167: show system certificate Output Fields	2353

About the Documentation

- Documentation and Release Notes on page lvii
- Supported Platforms on page lvii
- Using the Examples in This Manual on page lvii
- Documentation Conventions on page lix
- Documentation Feedback on page lxi
- Requesting Technical Support on page lxi

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

Supported Platforms

For the features described in this document, the following platforms are supported:

- M Series
- MX Series
- T Series

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xml; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see the *CLI User Guide*.

Documentation Conventions

Table 1 on page [lix](#) defines notice icons used in this guide.

Table 1: Notice Icons







Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page [lx](#) defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric metric>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (string1 string2 string3)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [community-ids]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page at the Juniper Networks Technical Documentation site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes:
<http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications:
<http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum:
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.

PART 1

Services Interfaces Overview

- [Overview on page 3](#)
- [Configuration Overview on page 9](#)

CHAPTER 1

Overview

- [Understanding Services PICs on page 3](#)
- [Multiservices MIC and Multiservices MPC \(MS-MIC and MS-MPC\) Overview on page 5](#)
- [Supported Platforms on page 7](#)

Understanding Services PICs

Interfaces used in router networks can be broadly classified into two:

- Networking interfaces, such as Ethernet and SONET interfaces, that primarily provide traffic connectivity. For more information on these interfaces, see the Junos[®] OS Network Interfaces.
- Services interfaces, such as Adaptive Services interfaces and Multiservices interfaces, that provide specific capabilities for manipulating traffic before it is delivered to its destination.

Services interfaces enable you to add services to your network incrementally. Junos OS supports the following services interfaces:

- [Adaptive services and Multiservices PICs on page 3](#)
- [Encryption Services \(ES\) PIC on page 4](#)
- [Multilink Services and Link Services PICs on page 4](#)
- [Monitoring Services PICs on page 4](#)
- [Tunnel Services PIC on page 5](#)
- [Multiservices MIC and Multiservices MPC on page 5](#)

Adaptive services and Multiservices PICs

Adaptive Services [AS] PICs and Multiservices PICs enable you to perform multiple services on the same PIC by configuring a set of services and applications. The AS and Multiservices PICs offer a range of services that you can configure in one or more service sets. The following are some of the services you can configure on Adaptive services or multiservices interfaces:

- Class-of-service
- Intrusion detection service (IDS)

- IP Security (IPsec)
- Layer 2 tunneling protocols
- Monitoring services
- Network Address Translation (NAT)
- Stateful firewalls
- Voice services

For more information about these services, see [“Adaptive Services Overview” on page 25](#).



NOTE: On Juniper Networks MX Series 3D Universal Edge Routers, the Multiservices DPC provides essentially the same capabilities as the Multiservices PIC. The interfaces on both platforms are configured in the same way.

Encryption Services (ES) PIC

ES PIC provides a security suite for the IP version 4 (IPv4) and IP version 6 (IPv6) network layers. The suite provides functionality such as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. It also defines mechanisms for key generation and exchange, management of security associations, and support for digital certificates. For more information about encryption interfaces, see [“Configuring Encryption Interfaces” on page 1251](#).

Multilink Services and Link Services PICs

Multilink Services and Link Services PICs enable you to split, recombine, and sequence datagrams across multiple logical data links. The goal of multilink operation is to coordinate multiple independent links between a fixed pair of systems, providing a virtual link with greater bandwidth than any of the members. The Junos OS supports two services PICs based on the Multilink Protocol: the Multilink Services PIC and the Link Services PIC.

For more information about multilink and link services interfaces, see *Link and Multilink Services Interfaces Feature Guide for Routing Devices*.

Monitoring Services PICs

Monitoring Services PICs enable you to monitor traffic flow and export the monitored traffic. Monitoring traffic allows you to perform the following tasks:

- Gather and export detailed information about IPv4 traffic flows between source and destination nodes in your network.
- Sample all incoming IPv4 traffic on the monitoring interface and present the data in cflowd record format.
- Perform discard accounting on an incoming traffic flow.

- Encrypt or tunnel outgoing cflowd records, intercepted IPv4 traffic, or both.
- Direct filtered traffic to different packet analyzers and present the data in its original format.

For more information about flow monitoring interfaces, see *Monitoring, Sampling, and Collection Services Interfaces Feature Guide for Routing Devices*.

Tunnel Services PIC

Tunnel Services PIC provides a private, secure path through an otherwise public network by encapsulating arbitrary packets inside a transport protocol. Tunnels connect discontinuous subnetworks and enable encryption interfaces, virtual private networks (VPNs), and MPLS.

For more information about tunnel interfaces, see *Tunnel Properties*.

Multiservices MIC and Multiservices MPC

The Multiservices Modular Interfaces Card (MS-MIC) and the Multiservices Modular PIC Concentrator (MS-MPC), introduced in Junos OS Release 13.2, provide improved scaling and high performance. The MS-MIC and MS-MPC have enhanced memory (16 GB for MS-MIC, 32 GB per NPU of MS-MPC) and processing capabilities.

The services interfaces on MS-MPC and MS-MIC are identified in the configuration with an **ms-** prefix (for example, **ms-1/2/1**).

The following services packages come preinstalled and preconfigured on MS-MICs and MS-MPCs in Junos OS Release 13.2:

- Junos Traffic Vision (formerly referred to as Jflow/Flow Monitoring)
- Junos Address Aware (formerly referred to as NAT features)
- Junos VPN Site Secure (formerly referred to as IPsec features)
- Junos Network Secure (formerly referred to as the Stateful Firewall feature)

For information about MS-MIC and MS-MPC, see “[Multiservices MIC and Multiservices MPC \(MS-MIC and MS-MPC\) Overview](#)” on page 5.

Related Documentation

- [Supported Platforms on page 7](#)
- [Packet Flow Through the Adaptive Services or Multiservices PIC on page 27](#)
- [Enabling Service Packages on page 11](#)
- [Services Configuration Procedure on page 15](#)
- [Services Interface Naming Overview on page 9](#)

Multiservices MIC and Multiservices MPC (MS-MIC and MS-MPC) Overview

Juniper Networks supports the Multiservices Modular Interfaces Card (MS-MIC) and the Multiservices Modular PIC Concentrator (MS-MPC) that provide improved scaling and

high performance. The MS-MIC and MS-MPC have enhanced memory (16 GB for MS-MIC, 32 GB per NPU of MS-MPC) and processing capabilities.

The services interfaces on MS-MPC and MS-MIC are identified in the configuration with an **ms-** prefix (for example, **ms-1/2/1**). The following services packages come preinstalled and preconfigured on MS-MICs and MS-MPCs:

- Junos Traffic Vision (formerly referred to as Jflow)
- Junos Address Aware (formerly referred to as NAT features)
- Junos VPN Site Secure (formerly referred to as IPsec features)
- Junos Network Secure (formerly referred to as the Stateful Firewall feature)
- Junos Services Crypto Base PIC Package
- Junos Services Application Level Gateways



NOTE: You can check the default packages on an MS-MIC or MS-MPC by executing the **show extension-provider system packages interface ms-interace** operational mode command.

The MS-MIC and MS-MPC also support the captive portal content delivery (HTTP redirect) service package when configured for installation using the **set chassis** operational mode command.

Table 3 on page 6 lists the platforms on which the MS-MIC and MS-MPC are supported.

Table 3: MX Series Routers that Support MS-MIC and MS-MPC

	MX5	MX10	MX40	MX80	MX240	MX480	MX960	MX2010	MX2020
MS-MIC	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
NOTE: Only Junos Traffic Vision is supported.									
MS-MPC	No	No	No	No	Yes	Yes	Yes	No	No

You can install an MS-MIC on one of the following line cards:

- MPC-Type1
- MPC-Type2

Related Documentation

- [Understanding Services PICs on page 3](#)
- [Example: Configuring Junos VPN Site Secure on MS MIC and MS-MPC](#)
- [Example: Configuring Flow Monitoring on MS-MIC and MS-MPC](#)

Supported Platforms

For information about which platforms support Adaptive Services and MultiServices PICs and their features, see [“Enabling Service Packages” on page 11](#).

For information about PIC support on a specific Juniper Networks M Series Multiservice Edge Router or T Series Core Router, see the appropriate *PIC Guide* for the platform.

For information about MS-DPC support on a specific MX Series router, see the appropriate *DPC Guide* for the platform.

For information about services supported on Juniper Networks SRX Series Services Gateways, see [Feature Explorer](#).

- Related Documentation**
- [Understanding Services PICs on page 3](#)
 - [Multiservices MIC and Multiservices MPC \(MS-MIC and MS-MPC\) Overview on page 5](#)

CHAPTER 2

Configuration Overview

- [Services Interface Naming Overview on page 9](#)
- [Enabling Service Packages on page 11](#)
- [Services Configuration Procedure on page 15](#)
- [Example: Service Interfaces Configuration on page 16](#)
- [Configuring Default Timeout Settings for Services Interfaces on page 19](#)
- [Configuring System Logging for Services Interfaces on page 20](#)

Services Interface Naming Overview

Each interface has an interface name, which specifies the media type, the slot the FPC is located in, the location on the FPC that the PIC is installed in, and the PIC port. The interface name uniquely identifies an individual network connector in the system. You use the interface name when configuring interfaces and when enabling various functions and properties, such as routing protocols, on individual interfaces. The system uses the interface name when displaying information about the interface, for example, in the **show interfaces** command.

The interface name is represented by a physical part, a logical part, and a channel part in the following format:

physical<:channel>.logical

The channel part of the name is optional for all interfaces except channelized DS3, E1, OC12, and STM1 interfaces.

The physical part of an interface name identifies the physical device, which corresponds to a single physical network connector. This part of the interface name has the following format:

type-fpc/pic/port

type is the media type, which identifies the network device. For service interfaces, it can be one of the following:

- **ams**—Aggregated multiservices (AMS) interface. An AMS interface is a bundle of services interfaces that can function as a single interface. An AMS interface is denoted as **amsN** in the configuration, where **N** is a unique number that identifies an AMS interface (for example, **ams0**). The member interfaces in an AMS interface are identified in the configuration with an **mams-** prefix (for example, **mams-1/2/0**).
- **cp**—Flow collector interface.
- **es**—Encryption interface.
- **gr**—Generic routing encapsulation tunnel interface.
- **gre**—This interface is internally generated and not configurable.
- **ip**—IP-over-IP encapsulation tunnel interface.
- **ipip**—This interface is internally generated and not configurable.
- **ls**—Link services interface.
- **lsq**—Link services intelligent queuing (IQ) interface; also used for voice services.
- **mams**—Member interface in an AMS interface.
- **ml**—Multilink interface.
- **mo**—Monitoring services interface. The logical interface **mo-fpc/pic/port.16383** is an internally generated, nonconfigurable interface for router control traffic.
- **ms**—Multiservices interfaces on multiservices modular interfaces card (MS-MIC) and multiservices modular port concentrators (MS-MPC).
- **mt**—Multicast tunnel interface. This interface is automatically generated, but you can configure properties on it if needed.
- **mtun**—This interface is internally generated and not configurable.
- **rlsq**—Redundancy LSQ interface.
- **rsp**—Redundancy adaptive services interface.
- **si**—Services inline interface, configured on MX3D Series routers only.
- **sp**—Adaptive services interface. The logical interface **sp-fpc/pic/port.16383** is an internally generated, nonconfigurable interface for router control traffic.
- **tap**—This interface is internally generated and not configurable.
- **vt**—Virtual loopback tunnel interface.

**Related
Documentation**

- [Understanding Services PICs on page 3](#)
- [Understanding Aggregated Multiservices Interfaces on page 599](#)
- [Examples: Configuring Services Interfaces on page 44](#)

Enabling Service Packages

For AS PICs, Multiservices PICs, Multiservices DPCs, and the internal Adaptive Services Module (ASM) in the M7i router, there are two service packages: Layer 2 and Layer 3. Both service packages are supported on all adaptive services interfaces, but you can enable only one service package per PIC, with the exception of a combined package supported on the ASM. On a single router, you can enable both service packages by installing two or more PICs on the platform.



NOTE: Graceful Routing Engine switchover (GRES) is automatically enabled on all services PICs and DPCs except the ES PIC. It is supported on all M Series, MX Series, and T Series routers except for TX Matrix routers. Layer 3 services should retain state after switchover, but Layer 2 services will restart. For IPsec services, Internet Key Exchange (IKE) negotiations are not stored and must be restarted after switchover. For more information about GRES, see the *Junos OS High Availability Library for Routing Devices*.

You enable service packages per PIC, not per port. For example, if you configure the Layer 2 service package, the entire PIC uses the configured package. To enable a service package, include the service-package statement at the `[edit chassis fpc slot-number pic pic-number adaptive-services]` hierarchy level, and specify `layer-2` or `layer-3`:

```
[edit chassis fpc slot-number pic pic-number adaptive-services]
service-package (layer-2 | layer-3);
```

To determine which package an AS PIC supports, issue the `show chassis hardware` command: if the PIC supports the Layer 2 package, it is listed as **Link Services II**, and if it supports the Layer 3 package, it is listed as **Adaptive Services II**. To determine which package a Multiservices PIC supports, issue the `show chassis pic fpc-slot slot-number pic-slot slot-number` command. The **Package** field displays the value **Layer-2** or **Layer-3**.



NOTE: The ASM has a default option (`layer-2-3`) that combines the features available in the Layer 2 and Layer 3 service packages.

After you commit a change in the service package, the PIC is taken offline and then brought back online immediately. You do not need to manually take the PIC offline and online.



NOTE: Changing the service package causes all state information associated with the previous service package to be lost. You should change the service package only when there is no active traffic going to the PIC.

The services supported in each package differ by PIC and platform type. [Table 4 on page 12](#) lists the services supported within each service package for each PIC and platform.

On the AS and Multiservices PICs, *link services* support includes Junos OS CoS components, LFI (FRF.12), MLFR end-to-end (FRF.15), MLFR UNI NNI (FRF.16), MLPPP (RFC 1990), and multiclass MLPPP. For more information, see “[Layer 2 Service Package Capabilities and Interfaces](#)” on page 14 and “[Layer 2 Service Package Capabilities and Interfaces](#)” on page 543.



NOTE: The AS PIC II for Layer 2 Service is dedicated to supporting the Layer 2 service package only.

For additional information about Layer 3 services, see the *Junos OS, Release 14.2*.

Table 4: AS and Multiservices PIC Services by Service Package, PIC, and Platform

Services	ASM	AS/AS2 PICs and Multiservices PICs	AS/AS2 and Multiservices PICs	AS2 and Multiservices PICs	AS2 and Multiservices PICs
Layer 2 Service Package (Only)	M7i	M7i, M10i, and M20	M40e and M120	M320, T320, and T640	TX Matrix
Link Services:					
• Link services	Yes	Yes	Yes	Yes	No
• Multiclass MLPPP	Yes	Yes	Yes	Yes	No
Voice Services:					
• CRTP and LFI	Yes	Yes	Yes	Yes	No
• CRTP and MLPPP	Yes	Yes	Yes	Yes	No
• CRTP over PPP (without MLPPP)	Yes	Yes	Yes	Yes	No
Layer 3 Service Package (Only)	M7i	M7i, M10i, and M20	M40e and M120	M320, T320, and T640	TX Matrix
Security Services:					
• CoS	Yes	Yes	Yes	Yes	No
• Intrusion detection system (IDS)	Yes	Yes	Yes	Yes	No
• IPsec	Yes	Yes	Yes	Yes	No
• NAT	Yes	Yes	Yes	Yes	No
• Stateful firewall	Yes	Yes	Yes	Yes	No
Accounting Services:					

Table 4: AS and Multiservices PIC Services by Service Package, PIC, and Platform (*continued*)

Services	ASM	AS/AS2 PICs and Multiservices PICs	AS/AS2 and Multiservices PICs	AS2 and Multiservices PICs	AS2 and Multiservices PICs
• Active monitoring	Yes	Yes	Yes	Yes	Yes
• Dynamic flow capture (Multiservices 400 PIC only)	No	No	No	Yes	No
• Flow-tap	Yes	Yes	Yes (M40e only)	Yes	No
• Passive monitoring (Multiservices 400 PIC only)	No	Yes	Yes (M40e only)	Yes	No
• Port mirroring	Yes	Yes	Yes	Yes	Yes
LNS Services:					
• L2TP LNS	Yes	Yes (M7i and M10i only)	Yes (M120 only)	No	No
Voice Services:					
• BGF	Yes	Yes	Yes	Yes	No
Layer 2 and Layer 3 Service Package (Common Features)	M7i	M7i, M10i, and M20	M40e and M120	M320, T320, and T640	TX Matrix
RPM Services:					
• RPM probe timestamping	Yes	Yes	Yes	Yes	No
Tunnel Services:					
• GRE (<i>gr-fpc/pic/port</i>)	Yes	Yes	Yes	Yes	Yes
• GRE fragmentation (<i>clear-dont-fragment-bit</i>)	Yes	Yes	Yes	No	No
• GRE key	Yes	Yes	Yes	Yes	No
• IP-IP tunnels (<i>ip-fpc/pic/port</i>)	Yes	Yes	Yes	Yes	Yes
• Logical tunnels (<i>lt-fpc/pic/port</i>)	No	No	No	No	No
• Multicast tunnels (<i>mt-fpc/pic/port</i>)	Yes	Yes	Yes	Yes	Yes
• PIM de-encapsulation (<i>pd-fpc/pic/port</i>)	Yes	Yes	Yes	Yes	Yes

Table 4: AS and Multiservices PIC Services by Service Package, PIC, and Platform (*continued*)

Services	ASM	AS/AS2 PICs and Multiservices PICs	AS/AS2 and Multiservices PICs	AS2 and Multiservices PICs	AS2 and Multiservices PICs
• PIM encapsulation (pe-fpc/pic/port)	Yes	Yes	Yes	Yes	Yes
• Virtual tunnels (vt-fpc/pic/port)	Yes	Yes	Yes	Yes	Yes

Layer 2 Service Package Capabilities and Interfaces

When you enable the Layer 2 service package, you can configure link services. On the AS and Multiservices PICs and the ASM, link services include support for the following:

- Junos CoS components—“[Layer 2 Service Package Capabilities and Interfaces](#)” on [page 543](#) describes how the Junos CoS components work on link services IQ (**lsq**) interfaces. For detailed information about Junos CoS components, see the *Class of Service Feature Guide for Routing Devices*.
- LFI on Frame Relay links using FRF.12 end-to-end fragmentation—The standard for FRF.12 is defined in the specification FRF.12, *Frame Relay Fragmentation Implementation Agreement*.
- LFI on MLPPP links.
- MLFR UNI NNI (FRF.16)—The standard for FRF.16 is defined in the specification FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*.
- MLPPP (RFC 1990)
- MLFR end-to-end (FRF.15)

For the LSQ interface on the AS and Multiservices PICs, the configuration syntax is almost the same as for Multilink and Link Services PICs. The primary difference is the use of the interface-type descriptor **lsq** instead of **ml** or **ls**. When you enable the Layer 2 service package, the following interfaces are automatically created:

```
gr-fpc/pic/port
ip-fpc/pic/port
lsq-fpc/pic/port
lsq-fpc/pic/port:0
...
lsq-fpc/pic/port:N
mt-fpc/pic/port
pd-fpc/pic/port
pe-fpc/pic/port
sp-fpc/pic/port
vt-fpc/pic/port
```

Interface types **gr**, **ip**, **mt**, **pd**, **pe**, and **vt** are standard tunnel interfaces that are available on the AS and Multiservices PICs whether you enable the Layer 2 or the Layer 3 service package. These tunnel interfaces function the same way for both service packages,

except that the Layer 2 service package does not support some tunnel functions, as shown in [Table 4 on page 12](#).

Interface type `lsq-fpc/pic/port` is the physical link services IQ (`lsq`) interface. Interface types `lsq-fpc/pic/port:0` through `lsq-fpc/pic/port:N` represent FRF.16 bundles. These interface types are created when you include the `mlfr-uni-nni-bundles` statement at the `[edit chassis fpc slot-number pic pic-number]` option. For more information, see “[Layer 2 Service Package Capabilities and Interfaces](#)” on page 543 and *Link and Multilink Services Interfaces Feature Guide for Routing Devices*.



NOTE: Interface type `sp` is created because it is needed by the Junos OS. For the Layer 2 service package, the `sp` interface is not configurable, but you should not disable it.

Related Documentation

- [Understanding Services PICs on page 3](#)
- [Adaptive Services Overview on page 25](#)
- [Supported Platforms on page 7](#)
- [Packet Flow Through the Adaptive Services or Multiservices PIC on page 27](#)
- [Services Configuration Procedure on page 15](#)

Services Configuration Procedure

You follow these general steps to configure services:

1. Define application objects by configuring statements at the `[edit applications]` hierarchy level.
2. Define service rules by configuring statements at the `[edit services (ids | ipsec-vpn | nat | stateful-firewall) rule]` hierarchy level.
3. Group the service rules by configuring the `rule-set` statement at the `[edit services (ids | ipsec-vpn | nat | stateful-firewall)]` hierarchy level.
4. Group service rule sets under a service-set definition by configuring the `service-set` statement at the `[edit services]` hierarchy level.
5. Apply the service set on an interface by including the `service-set` statement at the `[edit interfaces interface-name unit logical-unit-number family inet service (input | output)]` hierarchy level. Alternatively, you can configure logical interfaces as a next-hop destination by including the `next-hop-service` statement at the `[edit services service-set service-set-name]` hierarchy level.



NOTE: You can configure IDS, NAT, and stateful firewall service rules within the same service set. You must configure IPsec services in a separate service set, although you can apply both service sets to the same PIC.

- Related Documentation**
- [Understanding Services PICs on page 3](#)
 - [Enabling Service Packages on page 11](#)
 - [Supported Platforms on page 7](#)

Example: Service Interfaces Configuration

The following configuration includes all the items necessary to configure services on an interface.

```
[edit]
interfaces {
  fe-0/1/0 {
    unit 0 {
      family inet {
        service {
          input {
            service-set Firewall-Set;
          }
          output {
            service-set Firewall-Set;
          }
        }
        address 10.1.3.2/24;
      }
    }
  }
  fe-0/1/1 {
    unit 0 {
      family inet {
        filter {
          input Sample;
        }
        address 172.16.1.2/24;
      }
    }
  }
  sp-1/0/0 {
    unit 0 {
      family inet {
        address 172.16.1.3/24 {
        }
      }
    }
  }
}
forwarding-options {
  sampling {
    input {
      family inet {
        rate 1;
      }
    }
  }
}
```

```
output {
  cflowd 10.1.3.1 {
    port 2055;
    version 5;
  }
  flow-inactive-timeout 15;
  flow-active-timeout 60;
  interface sp-1/0/0 {
    engine-id 1;
    engine-type 136;
    source-address 10.1.3.2;
  }
}
}
}
firewall {
  filter Sample {
    term Sample {
      then {
        count Sample;
        sample;
        accept;
      }
    }
  }
}
}
services {
  stateful-firewall {
    rule Rule1 {
      match-direction input;
      term 1 {
        from {
          application-sets Applications;
        }
        then {
          accept;
        }
      }
    }
    term accept {
      then {
        accept;
      }
    }
  }
  rule Rule2 {
    match-direction output;
    term Local {
      from {
        source-address {
          10.1.3.2/32;
        }
      }
      then {
        accept;
      }
    }
  }
}
```

```
    }
  }
  ids {
    rule Attacks {
      match-direction output;
      term Match {
        from {
          application-sets Applications;
        }
        then {
          logging {
            syslog;
          }
        }
      }
    }
  }
}
nat {
  pool public {
    address-range low 172.16.2.1 high 172.16.2.32;
    port automatic;
  }
  rule Private-Public {
    match-direction input;
    term Translate {
      then {
        translated {
          source-pool public;
          translation-type source napt-44;
        }
      }
    }
  }
}
service-set Firewall-Set {
  stateful-firewall-rules Rule1;
  stateful-firewall-rules Rule2;
  nat-rules Private-Public;
  ids-rules Attacks;
  interface-service {
    service-interface sp-1/0/0;
  }
}
applications {
  application ICMP {
    application-protocol icmp;
  }
  application FTP {
    application-protocol ftp;
    destination-port ftp;
  }
  application-set Applications {
    application ICMP;
    application FTP;
  }
}
```

```
}
```

Configuring Default Timeout Settings for Services Interfaces

You can specify global default settings for certain timers that apply for the entire interface. There are three statements of this type:

- **inactivity-timeout**—Sets the inactivity timeout period for established flows, after which they are no longer valid.
- **open-timeout**—Sets the timeout period for Transmission Control Protocol (TCP) session establishment, for use with SYN-cookie defenses against network intrusion.
- **close-timeout**—Sets the timeout period for Transmission Control Protocol (TCP) session tear-down.

To configure a setting for the inactivity timeout period, include the **inactivity-timeout** statement at the **[edit interfaces *interface-name* services-options]** hierarchy level:

```
[edit interfaces interface-name services-options]
  inactivity-timeout seconds;
```

The default value is 30 seconds. The range of possible values is from 4 through 86,400 seconds. Any value you configure in the application protocol definition overrides the value specified here; for more information, see [“Configuring Application Protocol Properties” on page 303](#).

To configure a setting for the TCP session establishment timeout period, include the **open-timeout** statement at the **[edit interfaces *interface-name* services-options]** hierarchy level:

```
[edit interfaces interface-name services-options]
  open-timeout seconds;
```

The default value is 5 seconds. The range of possible values is from 4 through 224 seconds. Any value you configure in the intrusion detection service (IDS) definition overrides the value specified here; for more information, see [Intrusion Detection Properties](#).

To configure a setting for the TCP session teardown timeout period, include the **close-timeout** statement at the **[edit interfaces *interface-name* services-options]** hierarchy level:

```
[edit interfaces interface-name services-options]
  close-timeout seconds;
```

The default value is 1 second. The range of possible values is from 2 through 300 seconds.

Use of Keep-Alive Messages for Greater Control of TCP Inactivity Timeouts

Keep-alive messages are generated automatically to prevent TCP inactivity timeouts. The default number of keep-alive messages is 4. However, you can configure the number of keep-alive messages by entering the **tcp-tickles** statement at the **[edit interfaces *interface-name* service-options]** hierarchy level.

When timeout is generated for a bidirectional TCP flow, keep-alive packets are sent to reset the timer. If number of consecutive keep-alive packets sent in a flow reaches the default or configured limit, the conversation is deleted. There are several possible scenarios, depending on the setting of the **inactivity-timer** and the default or configured maximum number of keep-alive messages.

- If the configured value of keep-alive messages is zero and **inactivity-timeout** is NOT configured (in which case the default timeout value of 30 is used), no keep-alive packets are sent. The conversation is deleted when any flow in the conversation is idle for more than 30 seconds.
- If the configured value of keep-alive messages is zero and the **inactivity-timeout** is configured, no keep-alive packets are sent, and the conversation is deleted when any flow in the conversation is idle for more than the configured timeout value.
- If the default or configured maximum number of keep-alive messages is some positive integer, and any of the flows in a conversation is idle for more than the default or configured value for **inactivity-timeout** keep-alive packets are sent. If hosts do not respond to the configured number of consecutive keep-alive packets, the conversation is deleted. The interval between keep-alive packets will be 1 second. However, if the host sends back an ACK packet, the corresponding flow becomes active, and keep-alive packets are not sent until the flow becomes idle again.

**Related
Documentation**

- [Understanding Services PICs on page 3](#)
- [Configuring the Address and Domain for Services Interfaces on page 45](#)
- [Configuring System Logging for Services Interfaces on page 20](#)
- [Applying Filters and Services to Interfaces on page 38](#)
- [Examples: Configuring Services Interfaces on page 44](#)

Configuring System Logging for Services Interfaces

You specify properties that control how system log messages are generated for the interface as a whole. If you configure different values for the same properties at the **[edit services service-set service-set-name]** hierarchy level, the service-set values override the values configured for the interface. For more information on configuring service-set properties, see [“Configuring System Logging for Service Sets” on page 47](#).

To configure interface-wide default system logging values, include the **syslog** statement at the **[edit interfaces interface-name services-options]** hierarchy level:

```
[edit interfaces interface-name services-options]
syslog {
  host hostname {
    services severity-level;
    facility-override facility-name;
    log-prefix prefix-value;
    port port-number;
  }
}
```

Configure the **host** statement with a hostname or an IP address that specifies the system log target server. The hostname **local** directs system log messages to the Routing Engine. For external system log servers, the hostname must be reachable from the same routing instance to which the initial data packet (that triggered session establishment) is delivered. You can specify only one system logging hostname.

Table 5 on page 21 lists the severity levels that you can specify in configuration statements at the **[edit interfaces *interface-name* services-options syslog host *hostname*]** hierarchy level. The levels from **emergency** through **info** are in order from highest severity (greatest effect on functioning) to lowest.

Table 5: System Log Message Severity Levels

Severity Level	Description
any	Includes all severity levels
emergency	System panic or other condition that causes the router to stop functioning
alert	Conditions that require immediate correction, such as a corrupted system database
critical	Critical conditions, such as hard drive errors
error	Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels
warning	Conditions that warrant monitoring
notice	Conditions that are not errors but might warrant special handling
info	Events or nonerror conditions of interest

We recommend setting the system logging severity level to **error** during normal operation. To monitor PIC resource usage, set the level to **warning**. To gather information about an intrusion attack when an intrusion detection system error is detected, set the level to **notice** for a specific interface. To debug a configuration or log Network Address Translation (NAT) functionality, set the level to **info**.

For more information about system log messages, see the *Junos OS System Log Messages Reference*.

To use one particular facility code for all logging to the specified system log host, include the **facility-override** statement at the **[edit interfaces *interface-name* services-options syslog host *hostname*]** hierarchy level:

```
[edit interfaces interface-name services-options]
  facility-override facility-name;
```

The supported facilities include **authorization**, **daemon**, **ftp**, **kernel**, **user**, and **local0** through **local7**.

To specify a text prefix for all logging to this system log host, include the **log-prefix** statement at the **[edit interfaces *interface-name* services-options syslog host *hostname*]** hierarchy level:

```
[edit interfaces interface-name services-options]  
log-prefix prefix-value;
```

**Related
Documentation**

- [Understanding Services PICs on page 3](#)
- [Configuring the Address and Domain for Services Interfaces on page 45](#)
- [Configuring Default Timeout Settings for Services Interfaces on page 19](#)
- [Applying Filters and Services to Interfaces on page 38](#)
- [Examples: Configuring Services Interfaces on page 44](#)

PART 2

Adaptive Services Overview

- [Adaptive Services Overview on page 25](#)
- [Adaptive Services Configuration Overview on page 29](#)

CHAPTER 3

Adaptive Services Overview

- [Adaptive Services Overview on page 25](#)
- [Packet Flow Through the Adaptive Services or Multiservices PIC on page 27](#)

Adaptive Services Overview

MultiServices PICs and MultiServices Dense Port Concentrators (MS-DPCs) provide *adaptive services interfaces*, which allow you to coordinate multiple services on a single PIC by configuring a set of services and applications. MultiServices PICs and MS-DPCs offer a special range of services you configure in one or more service sets.

The MultiServices PIC is available in three versions, the MultiServices 100, the MultiServices 400, and the MultiServices 500, which differ in memory size and performance. All versions offer enhanced performance in comparison with AS PICs. MultiServices PICs are supported on M Series and T Series routers except M20 routers.

The MultiServices DPC is available for MX Series routers; it includes a subset of the functionality supported on the MultiServices PIC. Currently the MultiServices DPC supports the following Layer 3 services: stateful firewall, NAT, IDS, IPsec, active flow monitoring, RPM, and generic routing encapsulation (GRE) tunnels (including GRE key and fragmentation); it also supports graceful Routing Engine switchover (GRES) and Dynamic Application Awareness for Junos OS. For more information about supported packages, see [“Enabling Service Packages” on page 11](#).

It is also possible to group several Multiservices PICs into an aggregated Multiservices (AMS) system. An AMS configuration eliminates the need for separate routers within a system. The primary benefit of having an AMS configuration is the ability to support load balancing of traffic across multiple services PICs. Starting with Junos OS 11.4, all MX Series routers will support high availability (HA) and Network Address Translation (NAT) on AMS infrastructure. See [“Configuring Load Balancing on AMS Infrastructure” on page 605](#) for more information.



NOTE: The MultiServices PICs are polling based and not interrupt based; as a result, a high value in the `show chassis pic` “Interrupt load average” field may not mean that the PIC has reached its maximum limit of processing.

The following services are configured within a service set and are available only on adaptive services interfaces:

- Stateful firewall—A type of firewall filter that considers state information derived from previous communications and other applications when evaluating traffic.
- Network Address Translation (NAT)—A security procedure for concealing host addresses on a private network behind a pool of public addresses.
- Intrusion detection service (IDS)—A set of tools for detecting, redirecting, and preventing certain kinds of network attack and intrusion.
- IP Security (IPsec)—A set of tools for configuring manual or dynamic security associations (SAs) for encryption of data traffic.
- Class of service (CoS)—A subset of CoS functionality for services interfaces, limited to DiffServ code point (DSCP) marking and forwarding-class assignment. CoS BA classification is not supported on services interfaces.

The configuration for these services comprises a series of rules that you can arrange in order of precedence as a *rule set*. Each rule follows the structure of a firewall filter, with a **from** statement containing input or match conditions and a **then** statement containing actions to be taken if the match conditions are met.

The following services are also configured on the MultiServices PICs and MS-DPCs, but do not use the rule set definition:

- Layer 2 Tunneling Protocol (L2TP)—A tool for setting up secure tunnels using Point-to-Point Protocol (PPP) encapsulation across Layer 2 networks.
- Link Services Intelligent Queuing (LSQ)—Interfaces that support Junos OS class-of-service (CoS) components, link fragmentation and interleaving (LFI) (FRF.12), Multilink Frame Relay (MLFR) user-to-network interface (UNI) network-to-network interface (NNI) (FRF.16), and Multilink PPP (MLPPP).
- Voice services—A feature that uses the Compressed Real-Time Transport Protocol (CRTP) to enable voice over IP traffic to use low-speed links more effectively.

In addition, Junos OS includes the following tools for configuring services:

- Application protocols definition—Allows you to configure properties of application protocols that are subject to processing by router services, and group the application definitions into application sets.
- Service-set definition—Allows you to configure combinations of directional rules and default settings that control the behavior of each service in the service set.



NOTE: Logging of adaptive services interfaces messages to an external server by means of the `fxp0` port is not supported on M Series routers. The architecture does not support system logging traffic out of a management interface. Instead, access to an external server is supported on a Packet Forwarding Engine interface.

- Related Documentation**
- [Understanding Services PICs on page 3](#)
 - [Packet Flow Through the Adaptive Services or Multiservices PIC on page 27](#)
 - [Enabling Service Packages on page 11](#)
 - [Services Configuration Procedure on page 15](#)
 - [Supported Platforms on page 7](#)

Packet Flow Through the Adaptive Services or Multiservices PIC

You can optionally configure service sets to be applied at one of the following three points while the packets transit the router:

- An interface service set applied at the inbound interface.
- A next-hop service set applied at the forwarding table.
- An interface service set applied at the outbound interface.

The packet flow is as follows, graphically displayed in [Figure 1 on page 28](#). (You can configure a service set as either an interface service set or a next-hop service set.)

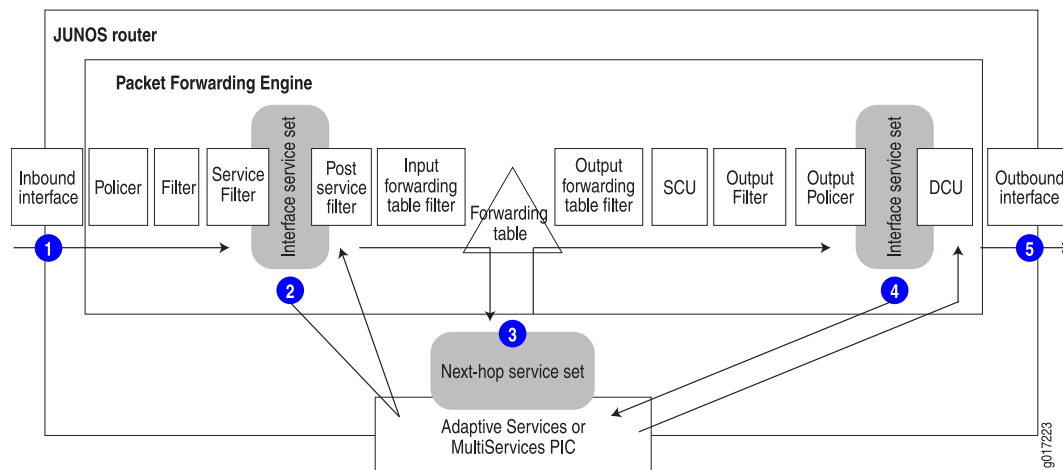
1. Packets enter the router on the inbound interface.
2. A policer, filter, service filter, service set, postservice filter, and input forwarding-table filter are applied sequentially to the traffic; these are all optional items in the configuration. If an interface service set is applied, the packets are forwarded to the AS or MultiServices PIC for services processing and then sent back to the Packet Forwarding Engine; if a service filter is also applied, only packets matching the service filter are sent to the PIC. The optional postservice filter is applied and postprocessing takes place.
3. A next-hop service set can be applied to the VPN routing and forwarding (VRF) table or to **inet.0**. If it is applied, packets are sent to the PIC for services processing and sent back to the Packet Forwarding Engine.



NOTE: For NAT, the next-hop service set can only be applied to the VRF table. For all other services, the next-hop service set can be applied to either the VRF table or to **inet.0**.

4. On the output interface, an output filter, output policer, and interface service set can be applied sequentially to the traffic if you have configured any of these items. If an interface service set is applied, the traffic is forwarded to the PIC for processing and sent back to the Packet Forwarding Engine, which then forwards the traffic.
5. Packets exit the router.

Figure 1: Packet Flow Through the Adaptive Services or MultiServices PIC



NOTE: When an AS PIC experiences persistent back pressure as a result of high traffic volume for 3 seconds, the condition triggers an automatic core dump and reboot of the PIC to help clear the blockage. A system log message at level LOG_ERR is generated. This mechanism applies to both Layer 2 and Layer 3 service packages.

Related Documentation

- [Understanding Services PICs on page 3](#)
- [Adaptive Services Overview on page 25](#)
- [Supported Platforms on page 7](#)
- [Services Configuration Procedure on page 15](#)

CHAPTER 4

Adaptive Services Configuration Overview

- [Understanding Service Sets on page 29](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 31](#)
- [Configuring Service Rules on page 36](#)
- [Configuring Service Set Limitations on page 37](#)
- [Enabling Services PICs to Accept Multicast Traffic on page 38](#)
- [Applying Filters and Services to Interfaces on page 38](#)
- [Example: Configuring Service Sets on page 41](#)
- [Configuring AS or Multiservices PIC Redundancy on page 41](#)
- [Examples: Configuring Services Interfaces on page 44](#)
- [Configuring the Address and Domain for Services Interfaces on page 45](#)
- [Configuring System Logging for Service Sets on page 47](#)
- [Tracing Services PIC Operations on page 48](#)

Understanding Service Sets

Junos OS enables you to create service sets that define a collection of services to be performed by an Adaptive Services interface (AS) or Multiservices line cards (MS-DPC, MS-MIC, and MS-MPC). You can configure the service set either as an interface style service set or as a next-hop style service set.

An interface service set is used as an action modifier across an entire interface. You can use an interface style service set when you want to apply services to packets passing through an interface.

A next-hop service set is a route-based method of applying a particular service. Only packets destined for a specific next hop are serviced by the creation of explicit static routes. This configuration is useful when services need to be applied to an entire virtual private network (VPN) routing and forwarding (VRF) table, or when routing decisions determine that services need to be performed. When a next-hop service is configured, the service interface is considered to be a two-legged module with one leg configured to be the inside interface (inside the network) and the other configured as the outside interface (outside the network).

To configure service sets, include the following statements at the **[edit services]** hierarchy level:

```
[edit services]
service-set service-set-name {
  (ids-rules rule-names | ids-rule-sets rule-set-name);
  (ipsec-vpn-rules rule-names | ipsec-vpn-rule-sets rule-set-name);
  (nat-rules rule-names | nat-rule-sets rule-set-name);
  (pgcp-rules rule-names | pgcp-rule-sets rule-set-name);
  (ptsp-rules rule-names | ptsp-rule-sets rule-set-name);
  (stateful-firewall-rules rule-names | stateful-firewall-rule-sets rule-set-name);
  allow-multicast;
  extension-service service-name {
    provider-specific rules;
  }
  interface-service {
    service-interface interface-name;
  }
  ipsec-vpn-options {
    anti-replay-window-size bits;
    clear-dont-fragment-bit;
    ike-access-profile profile-name;
    local-gateway address;
    no-anti-replay;
    passive-mode-tunneling;
    trusted-ca [ ca-profile-names ];
    tunnel-mtu bytes;
  }
  max-flows number;
  next-hop-service {
    inside-service-interface interface-name.unit-number;
    outside-service-interface interface-name.unit-number;
    service-interface-pool name;
  }
  syslog {
    host hostname {
      services severity-level;
      facility-override facility-name;
      log-prefix prefix-value;
    }
  }
}
adaptive-services-pics {
  traceoptions {
    file filename <files number> <match regex> <size size> <(world-readable |
      no-world-readable)>;
    flag flag;
  }
}
logging {
  traceoptions {
    file filename <files number> <match regex> <size size> <(world-readable |
      no-world-readable)>;
    flag flag;
  }
}
```


}

Related Documentation

- [Configuring Service Sets to be Applied to Services Interfaces on page 31](#)
- [Configuring Service Rules on page 36](#)
- [Configuring IPsec Service Sets on page 430](#)
- [Configuring Service Set Limitations on page 37](#)
- [Configuring System Logging for Service Sets on page 47](#)
- [Enabling Services PICs to Accept Multicast Traffic on page 38](#)
- [Tracing Services PIC Operations on page 48](#)
- [Example: Configuring Service Sets on page 41](#)

Configuring Service Sets to be Applied to Services Interfaces

You configure a services interface to specify the adaptive services interface on which the service is to be performed. Services interfaces are used with either of the service set types described in the following sections.

- [Configuring Interface Service Sets on page 31](#)
- [Configuring Next-Hop Service Sets on page 33](#)
- [Determining Traffic Direction on page 34](#)

Configuring Interface Service Sets

An interface service set is used as an action modifier across an entire interface. To configure the services interface, include the **interface-service** statement at the **[edit services service-set service-set-name]** hierarchy level:

```
[edit services service-set service-set-name]
interface-service {
  service-interface interface-name;
}
```

Only the device name is needed, because the router software manages logical unit numbers automatically. The services interface must be an adaptive services interface for which you have configured **unit 0 family inet** at the **[edit interfaces interface-name]** hierarchy level.

When you have defined and grouped the service rules by configuring the service-set definition, you can apply services to one or more interfaces installed on the router. When you apply the service set to an interface, it automatically ensures that packets are directed to the PIC.

To associate a defined service set with an interface, include a **service-set** statement with the **input** or **output** statement at the **[edit interfaces interface-name unit logical-unit-number family inet service]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet service]
input {
```

```
service-set service-set-name <service-filter filter-name>;
post-service-filter filter-name;
}
output {
  service-set service-set-name <service-filter filter-name>;
}
```

If a packet is entering the interface, the match direction is **input**. If a packet is leaving the interface, the match direction is **output**. The service set retains the input interface information even after services are applied, so that functions such as filter-class forwarding and destination class usage (DCU) that depend on input interface information continue to work.

You configure the same service set on the input and output sides of the interface. You can optionally include filters associated with each service set to refine the target and additionally process the traffic. If you include the **service-set** statement without a **service-filter** definition, the router software assumes the match condition is true and selects the service set for processing automatically.



NOTE: If you configure service sets with filters, they must be configured on the input and output sides of the interface.

You can include more than one service set definition on each side of the interface. If you include multiple service sets, the router software evaluates them in the order in which they appear in the configuration. The system executes the first service set for which it finds a match in the service filter and ignores the subsequent definitions. A maximum of six service sets can be applied to an interface. When you apply multiple service sets to an interface, you must also configure and apply a service filter to the interface.

An additional statement allows you to specify a filter for processing the traffic after the input service set is executed. To configure this type of filter, include the **post-service-filter** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet service input]** hierarchy level:

```
post-service-filter filter-name;
```

The **post-service-filter** statement is not supported when the service interface is on an MS-MIC or MS-MPC.

For an example, see [“Example: Configuring Service Sets” on page 41](#).



NOTE: With interface-style service sets that are configured with Junos OS extension-provide packages, the traffic fails to get serviced when the ingress interface is part of a VRF instance and the service interface is not part of the same VRF instance.



NOTE: When the MultiServices PIC configured for a service set is either administratively taken offline or undergoes a failure, all the traffic entering the configured interface with an IDP service set would be dropped without notification. To avoid this traffic loss, include the `bypass-traffic-on-pic-failure` statement at the `[edit services service-set service-set-name service-set-options]` hierarchy level. When this statement is configured, the affected packets are forwarded in the event of a MultiServices PIC failure or offlining, as though interface-style services were not configured. This issue applies only to Junos Application Aware (previously known as Dynamic Application Awareness) configurations using IDP service sets. This forwarding feature worked only with the Packet Forwarding Engine (PFE) initially. Starting with Junos OS Release 11.3, the packet-forwarding feature is extended to packets generated by the Routing Engine for bypass service sets as well.

Configuring Next-Hop Service Sets

A next-hop service set is a route-based method of applying a particular service. Only packets destined for a specific next hop are serviced by the creation of explicit static routes. This configuration is useful when services need to be applied to an entire virtual private network (VPN) routing and forwarding (VRF) table, or when routing decisions determine that services need to be performed.

When a next-hop service is configured, the AS or Multiservices PIC is considered to be a two-legged module with one leg configured to be the inside interface (inside the network) and the other configured as the outside interface (outside the network).



NOTE: You can create IFL indexes greater than 8000 only if the interface service set is not configured.

To configure the domain, include the `service-domain` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level:

```
service-domain (inside | outside);
```

The `service-domain` setting must match the configuration for the next-hop service inside and outside interfaces. To configure the inside and outside interfaces, include the `next-hop-service` statement at the `[edit services service-set service-set-name]` hierarchy level. The interfaces you specify must be logical interfaces on the same AS PIC. You cannot configure `unit 0` for this purpose, and the logical interface you choose must not be used by another service set.

```
next-hop-service {
  inside-service-interface interface-name.unit-number;
  outside-service-interface interface-name.unit-number;
}
```

Traffic on which the service is applied is forced to the inside interface using a static route. For example:

```
routing-options {  
  static {  
    route 10.1.2.3 next-hop sp-1/1/0.1;  
  }  
}
```

After the service is applied, traffic exits by way of the outside interface. A lookup is then performed in the Packet Forwarding Engine (PFE) to send the packet out of the AS or Multiservices PIC.

The reverse traffic enters the outside interface, is serviced, and sent to the inside interface. The inside interface forwards the traffic out of the AS or Multiservices PIC.

Determining Traffic Direction

When you configure next-hop service sets, the AS PIC functions as a two-part interface, in which one part is the *inside* interface and the other part is the *outside* interface. The following sequence of actions takes place:

1. To associate the two parts with logical interfaces, you configure two logical interfaces with the **service-domain** statement, one with the **inside** value and one with the **outside** value, to mark them as either an inside or outside service interface.
2. The router forwards the traffic to be serviced to the inside interface, using the next-hop lookup table.
3. After the service is applied, the traffic exits from the outside interface. A route lookup is then performed on the packets to be sent out of the router.
4. When the reverse traffic returns on the outside interface, the applied service is undone; for example, IPsec traffic is decrypted or NAT addresses are unmasked. The serviced packets then emerge on the inside interface, the router performs a route lookup, and the traffic exits the router.

A service rule's match direction, whether input, output, or input/output, is applied with respect to the traffic flow through the AS PIC, not through a specific inside or outside interface.

When a packet is sent to an AS PIC, packet direction information is carried along with it. This is true for both interface style and next-hop style service sets.

Interface Style Service Sets

Packet direction is determined by whether a packet is entering or leaving any Packet Forwarding Engine interface (with respect to the forwarding plane) on which the **interface-service** statement is applied. This is similar to the input and output direction for stateless firewall filters.

The match direction can also depend on the network topology. For example, you might route all the external traffic through one interface that is used to protect the other interfaces on the router, and configure various services on this interface specifically.

Alternatively, you might use one interface for priority traffic and configure special services on it, but not care about protecting traffic on the other interfaces.

Next-Hop Style Service Sets

Packet direction is determined by the AS PIC interface used to route packets to the AS PIC. If you use the **inside-interface** statement to route traffic, then the packet direction is **input**. If you use the **outside-interface** statement to direct packets to the AS PIC, then the packet direction is **output**.

The interface to which you apply the service sets affects the match direction. For example, apply the following configuration:

```
sp-1/1/0 unit 1 service-domain inside;
sp-1/1/0 unit 2 service-domain outside;
```

If you configure **match-direction input**, you include the following statements:

```
[edit]
services service-set test1 next-hop-service inside-service-interface sp-1/0/0.1;
services service-set test1 next-hop-service outside-service-interface sp-1/0/0.2;
services ipsec-vpn rule test-ipsec-rule match-direction input;
routing-options static route 10.0.0.0/24 next-hop sp-1/1/0.1;
```

If you configure **match-direction output**, you include the following statements:

```
[edit]
services service-set test2 next-hop-service inside-service-interface sp-1/0/0.1;
services service-set test2 next-hop-service outside-service-interface sp-1/0/0.2;
services ipsec-vpn rule test-ipsec-rule match-direction output;
routing-options static route 10.0.0.0/24 next-hop sp-1/1/0.2;
```

The essential difference between the two configurations is the change in the match direction and the static routes' next hop, pointing to either the AS PIC's inside or outside interface.

Related Documentation

- [Understanding Service Sets on page 29](#)
- [Configuring Service Rules on page 36](#)
- [Configuring IPsec Service Sets on page 430](#)
- [Configuring Service Set Limitations on page 37](#)
- [Configuring System Logging for Service Sets on page 47](#)
- [Example: Configuring Service Sets on page 41](#)

Configuring Service Rules

You specify the collection of rules and rule sets that constitute the service set. The router performs rule sets in the order in which they appear in the configuration. You can include only one rule set for each service type. You configure the rule names and content for each service type at the **[edit services name]** hierarchy level for each type:

- You configure intrusion detection service (IDS) rules at the **[edit services ids]** hierarchy level; for more information, see [“Configuring IDS Rules” on page 355](#).
- You configure IP Security (IPsec) rules at the **[edit services ipsec-vpn]** hierarchy level; for more information, see *Junos VPN Site Secure*.
- You configure Network Address Translation (NAT) rules at the **[edit services nat]** hierarchy level; for more information, see *Junos Address Aware Carrier Grade NAT and IPv6 Feature Guide*.
- You configure packet-triggered subscribers and policy control (PTSP) rules at the **[edit services ptsp]** hierarchy level; for more information, see *Packet-Triggered Subscribers and Policy Control Feature Guide*.
- You configure softwire rules for DS-Lite or 6rd softwires at the **[edit services softwire]** hierarchy level; for more information, see *Softwire Services*.
- You configure stateful firewall rules at the **[edit services stateful-firewall]** hierarchy level; for more information, see *Junos Network Secure*.

To configure the rules and rule sets that constitute a service set, include the following statements at the **[edit services service-set service-set-name]** hierarchy level:

```
([ ids-rules rule-names ] | ids-rule-sets rule-set-name);
([ ipsec-vpn-rules rule-names ] | ipsec-vpn-rule-sets rule-set-name);
([ nat-rules rule-names ] | nat-rule-sets rule-set-name);
([ pgcp-rules rule-names ] | pgcp-rule-sets rule-set-name);
([ softwire-rules rule-names ] | softwire-rule-sets rule-set-name);
([ stateful-firewall-rules rule-names ] | stateful-firewall-rule-sets rule-set-name);
```

For each service type, you can include one or more individual rules, or one rule set.

If you configure a service set with IPsec rules, it must not contain rules for any other services. You can, however, configure another service set containing rules for the other services and apply both service sets to the same interface.



NOTE: You can also include Junos Application Aware (previously known as Dynamic Application Awareness) functionality within service sets. To do this, you must include an **idp-profile** statement at the **[edit services service-set]** hierarchy level, along with application identification (APPID) rules, and, as appropriate, application-aware access list (AACL) rules and a **policy-decision-statistics-profile**. Only one service sets can be applied to a single interface when Junos Application Aware functionality is used. For more information, see *Intrusion Detection and Prevention, Application Identification, and Application Aware Services Interfaces Feature Guide for Routing Devices*.

Related Documentation

- [Understanding Service Sets on page 29](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 31](#)
- [Configuring Service Set Limitations on page 37](#)
- [Configuring System Logging for Service Sets on page 47](#)

Configuring Service Set Limitations

You can set the following limitations on service set capacity:

- You can limit the maximum number of flows allowed per service set. To configure the maximum value, include the **max-flows** statement at the **[edit services service-set service-set-name]** hierarchy level:

max-flows *number*;

The **max-flows** statement permits you to assign a single flow limit value. For IDS service sets only, you can specify various types of flow limits with a finer degree of control. For more information, see the description of the **session-limit** statement in “[Configuring IDS Rule Sets](#)” on page 363.



NOTE: When an aggregated multiservices (AMS) interface is configured as the service interface for a service set, the **max-flow** value configured for the service set is applied to each of the member interfaces in the AMS interface. That is, if you have configured 1000 as the **max-flow** value for a service set that uses an AMS interface with four active member interfaces, each of the member interfaces can handle 1000 flows each, resulting in an effective **max-flow** value of 4000.

- You can limit the maximum segment size (MSS) allowed by the Transmission Control Protocol (TCP). To configure the maximum value, include the **tcp-mss** statement at the **[edit services service-set service-set-name]** hierarchy level:

tcp-mss *number*;

The TCP protocol negotiates an MSS value during session connection establishment between two peers. The MSS value negotiated is primarily based on the MTU of the interfaces to which the communicating peers are directly connected to. However in the network, due to variation in link MTU on the path taken by the TCP packets, some packets which are still well within the MSS value may be fragmented when the concerned packet's size exceeds the link's MTU.

If the router receives a TCP packet with the SYN bit and MSS option set and the MSS option specified in the packet is larger than the MSS value specified by the **tcp-mss** statement, the router replaces the MSS value in the packet with the lower value specified by the **tcp-mss** statement. The range for the **tcp-mss mss-value** parameter is from **536** through **65535**.

To view statistics of SYN packets received and SYN packets whose MSS value, is modified, issue the **show services service-sets statistics tcp-mss** operational mode

command. For more information on this topic, see the *Junos OS Administration Library for Routing Devices*.

- Related Documentation**
- [Understanding Service Sets on page 29](#)
 - [Configuring Service Sets to be Applied to Services Interfaces on page 31](#)
 - [Configuring Service Rules on page 36](#)
 - [Configuring System Logging for Service Sets on page 47](#)
 - [Configuring SNMP Traps for Flow Limits](#)

Enabling Services PICs to Accept Multicast Traffic

To allow multicast traffic to be sent to the Adaptive Services or Multiservices PIC, include the **allow-multicast** statement at the **[edit services service-set service-set-name]** hierarchy level. If this statement is not included, multicast traffic is dropped by default. This statement applies only to multicast traffic using a next-hop service set; interface service set configuration is not supported. Only unidirectional flows are created for multicast packets.

- Related Documentation**
- [Understanding Service Sets on page 29](#)
 - [Configuring Service Sets to be Applied to Services Interfaces on page 31](#)
 - [Configuring Service Rules on page 36](#)
 - [Example: Configuring Service Sets on page 41](#)
 - [Example: Configuring NAT for Multicast Traffic on page 105](#)

Applying Filters and Services to Interfaces

When you have defined and grouped the service rules by configuring the service-set definition, you can apply services to one or more interfaces on the router. To associate a defined service set with an interface, include the **service-set** statement with the **input** or **output** statement at the **[edit interfaces interface-name unit logical-unit-number family inet service]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet service]
input {
  service-set service-set-name <service-filter filter-name>;
  post-service-filter filter-name;
}
output {
  service-set service-set-name <service-filter filter-name>;
}
```



NOTE: When you enable services on an interface, reverse-path forwarding is not supported. You cannot configure services on the management interface (fxp0) or the loopback interface (lo0).

You can configure different service sets on the input and output sides of the interface. However, for service sets with bidirectional service rules, you must include the same service set definition in both the **input** and **output** statements. Any service set you include in the **service** statement must be configured with the **interface-service** statement at the **[edit services service-set *service-set-name*]** hierarchy level; for more information, see [“Configuring Service Sets to be Applied to Services Interfaces” on page 31](#).



NOTE: If you configure an interface with an input firewall filter that includes a reject action and with a service set that includes stateful firewall rules, the router executes the input firewall filter before the stateful firewall rules are run on the packet. As a result, when the Packet Forwarding Engine sends an Internet Control Message Protocol (ICMP) error message out through the interface, the stateful firewall rules might drop the packet because it was not seen in the input direction.

Possible workarounds are to include a forwarding-table filter to perform the reject action, because this type of filter is executed after the stateful firewall in the input direction, or to include an output service filter to prevent the locally generated ICMP packets from going to the stateful firewall service.

Configuring Service Filters

You can optionally include filters associated with each service set to refine the target and additionally process the traffic. If you include the **service-set** statement without a **service-filter** definition, the router software assumes that the match condition is true and selects the service set for processing automatically.

To configure service filters, include the **firewall** statement at the **[edit]** hierarchy level:

```
firewall {
  family inet {
    service-filter filter-name {
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
}
```



NOTE: You must specify **inet** as the address family to configure a service filter.

You configure service filters in a similar way to firewall filters. Service filters have the same match conditions as firewall filters, but the following specific actions:

- **count**—Add the packet to a counter total.
- **log**—Log the packet.
- **port-mirror**—Port-mirror the packet.
- **sample**—Sample the packet.
- **service**—Forward the packet for service processing.
- **skip**—Omit the packet from service processing.

For more information about configuring firewall filters, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

You can also include more than one service set definition on each side of the interface. If you include multiple service sets, the router software evaluates them in the order specified in the configuration. It executes the first service set for which it finds a match in the service filter and ignores the subsequent definitions.

An additional statement allows you to specify a filter for processing the traffic after the input service set is executed. To configure this type of filter, include the **post-service-filter** statement at the `[edit interfaces interface-name unit logical-unit-number family inet service input]` hierarchy level:

post-service-filter *filter-name*;



NOTE: The software performs postservice filtering only when it has selected and executed a service set. If the traffic does not meet the match criteria for any of the configured service sets, the postservice filter is ignored. The **post-service-filter** statement is not supported when the service interface is on an MS-MIC or MS-MPC.

For an example of applying a service set to an interface, see [“Examples: Configuring Services Interfaces” on page 44](#).

For more information on applying filters to interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*. For general information on filters, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.



NOTE: After NAT processing is applied to packets, they are not subject to output service filters. The service filters affect only untranslated traffic.

Related Documentation

- [Understanding Services PICs on page 3](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 31](#)
- [Examples: Configuring Services Interfaces on page 44](#)

Example: Configuring Service Sets

Apply two service sets, **my-input-service-set** and **my-output-service-set**, on an interface-wide basis. All traffic has **my-input-service-set** applied to it. After the service set is applied, additional filtering is done using **my_post_service_input_filter**.

```
[edit interfaces fe-0/1/0]
unit 0 {
  family inet {
    service {
      input {
        service-set my-input-service-set;
        post-service-filter my_post_service_input_filter;
      }
      output {
        service-set my-output-service-set;
      }
    }
  }
}
```

- Related Documentation**
- [Understanding Service Sets on page 29](#)
 - [Configuring Service Sets to be Applied to Services Interfaces on page 31](#)

Configuring AS or Multiservices PIC Redundancy

You can configure AS or Multiservices PIC redundancy on M Series and T Series routers, except TX Matrix routers, that have multiple AS or Multiservices PICs. To configure redundancy, you specify a redundancy services PIC (**rsp**) interface in which the primary PIC is active and a secondary PIC is on standby. If the primary PIC fails, the secondary PIC becomes active, and all service processing is transferred to it. If the primary AS or Multiservices PIC is restored, it remains on standby and does not preempt the secondary PIC; you need to manually restore the services to the primary PIC. To determine which PIC is currently active, issue the **show interfaces redundancy** command.

Failover to the secondary PIC occurs under the following conditions:

- The primary PIC, FPC, or Packet Forwarding Engine goes down, resets, or is physically removed from the router.
- The PIC or FPC is taken offline using the **request chassis pic fpc-slot slot-number pic-slot slot-number offline** or **request chassis fpc slot slot-number offline** command. For more information, see the [CLI Explorer](#).
- The driver watchdog timer expires.
- The **request interface switchover** command is issued. For more information, see the [CLI Explorer](#).



NOTE: Adaptive Services and Multiservices PICs in Layer-2 mode (running Layer 2 services) are not rebooted when a MAC flow-control situation is detected.



NOTE: When you perform a switchover from a primary PIC to a secondary or standby PIC or a revert operation by issuing request interfaces (`revert | switchover`) command for redundancy services PICs (`rsp`), the PIC that was previously the active PIC before the switchover or reversion is automatically rebooted. The reboot of the PIC that was previously active and functioning as the primary PIC does not disrupt traffic forwarding.

The physical interface type `rsp` specifies the pairings between primary and secondary `sp` interfaces to enable redundancy. To configure an AS or Multiservices PIC as the backup, include the `redundancy-options` statement at the `[edit interfaces rspnumber]` hierarchy level:

```
[edit interfaces rspnumber]
redundancy-options {
  primary sp-fpc/pic/port;
  secondary sp-fpc/pic/port;
  hot-standby;
}
```

For the `rsp` interface, *number* can be from 0 through 15.



NOTE: You can include a similar redundancy configuration for Link Services IQ (LSQ) PICs at the `[edit interfaces rlsqnumber]` hierarchy level. For more information, see “[Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces](#)” on page 548.

The following constraints apply to redundant AS or Multiservices PIC configurations:

- The services supported in redundancy configurations include stateful firewall, NAT, IDS, and IPsec. Services mounted on the AS or Multiservices PIC that use interface types other than **sp**- interfaces, such as tunneling and voice services, are not supported. For information on flow monitoring redundancy, see [“Configuring Services Interface Redundancy with Flow Monitoring” on page 826](#).



NOTE: For IPsec functionality, the router no longer needs to renegotiate security associations (SAs) during warm standby PIC switchover. Instead, the warm standby feature has been made stateful by periodically setting a checkpoint between the working state of the PIC and the Routing Engine, which should lessen the downtime during switchover. If you prefer to retain the earlier behavior, you can include the `clear-ipsec-sas-on-pic-restart` statement at the `[edit services ipsec-vpn]` hierarchy level. If you enable this capability, the router renegotiates the IPsec SAs on warm standby PIC switchover. For more information, see *Clearing Security Associations*.

- We recommend that you pair the same model type in RSP configurations, such as two ASMs or two AS2 PICs. If you pair unlike models, the two PICs may perform differently.
- You can specify an AS or Multiservices PIC (**sp** interface) as the primary for only one **rsp** interface.
- An **sp** interface can be a secondary for multiple **rsp** interfaces. However, the same **sp** interface cannot be configured as a primary interface in one **rsp** configuration and as a secondary in another configuration.
- When the secondary PIC is active, if another primary PIC that is paired with it in an **rsp** configuration fails, no failover takes place.
- When you configure an AS or Multiservices PIC within a redundant configuration, the **sp** interface cannot have any configured services. Apply the configurations at the `[edit interfaces rspnumber]` hierarchy level, using, for example, the **unit** and **services-options** statements. Exceptions include the **multiservice-options** statement used in flow monitoring configurations, which can be configured separately for the primary and secondary **sp** interfaces, and the **traceoptions** statement.
- All the operational mode commands that apply to **sp** interfaces also apply to **rsp** interfaces. You can issue **show** commands for the **rsp** interface or the primary and secondary **sp** interfaces.
- If a secondary PIC fails while it is in use, the **rsp** interface returns to the “not present” state. If the primary PIC comes up later, service is restored to it.
- For redundant Multiservices (**rms**-) interfaces, similar to the configuration of other bundle interfaces, the properties of the Multiservices (**ms**-) member interfaces, such as the logical unit and the address family, are inherited from the underlying **rms**- interface. If you previously configured the member **ms**- interface properties separately, and attempt to configure the **rms**- interface properties by using the relevant statements at the `[edit interfaces rmsnumber]` hierarchy level, an error occurs when you perform a commit check operation. You must configure the properties of interfaces that are part

of the `rms-` interface only by using the statements at the `[edit interfaces rmsnumber]` hierarchy level.

- Related Documentation**
- [Understanding Services PICs on page 3](#)
 - [Examples: Configuring Services Interfaces on page 44](#)
 - [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 608](#)

Examples: Configuring Services Interfaces

Apply the `my-service-set` service set on an interface-wide basis. All traffic that is accepted by `my_input_filter` has `my-input-service-set` applied to it. After the service set is applied, additional filtering is done using the `my_post_service_input_filter` filter.

```
[edit interfaces fe-0/1/0]
unit 0 {
  family inet {
    filter {
      input my_input_filter;
      output my_output_filter;
    }
    service {
      input {
        service-set my-input-service-set;
        post-service-filter my_post_service_input_filter;
      }
      output {
        service-set my-output-service-set;
      }
    }
  }
}
```

Configure two redundancy interfaces, `rsp0` and `rsp1`, and associated services.

```
[edit interfaces]
rsp0 {
  redundancy-options {
    primary sp-0/0/0;
    secondary sp-1/3/0;
  }
  unit 0 {
    family inet;
  }
  unit 30 {
    family inet;
    service-domain inside;
  }
  unit 31 {
    family inet;
    service-domain outside;
  }
}
rsp1 {
```

```

    redundancy-options {
        primary sp-0/1/0;
        secondary sp-1/3/0;
    }
    unit 0 {
        family inet;
    }
    unit 20 {
        family inet;
        service-domain inside;
    }
    unit 21 {
        family inet;
        service-domain outside;
    }
}
[edit services]
service-set null-sfw-with-nat {
    stateful-firewall-rules allow-all;
    nat-rules rule1;
    next-hop-service {
        inside-service-interface rsp0.30;
        outside-service-interface rsp0.31;
    }
}
[edit routing-instances]
vpna {
    interface rsp0.0;
}

```

Related Documentation

- [Understanding Services PICs on page 3](#)
- [Configuring the Address and Domain for Services Interfaces on page 45](#)
- [Configuring Default Timeout Settings for Services Interfaces on page 19](#)
- [Configuring System Logging for Services Interfaces on page 20](#)
- [Applying Filters and Services to Interfaces on page 38](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 608](#)

Configuring the Address and Domain for Services Interfaces

On the AS or Multiservices PIC, you configure a source address for system log messages by including the **address** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level:

```

address address {
    ...
}

```

Assign an IP address to the interface by configuring the **address** value. The AS or Multiservices PIC generally supports only IP version 4 (IPv4) addresses configured using the **family inet** statement, but IPsec services support IP version 6 (IPv6) addresses as well, configured using the **family inet6** statement.



NOTE: If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration, the remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
}
```

For more information on configuring the same address on multiple interfaces, see *Configuring the Interface Address*.

For information on other addressing properties you can configure that are not specific to service interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

The **service-domain** statement specifies whether the interface is used within the network or to communicate with remote devices. The software uses this setting to determine which default stateful firewall rules to apply, and to determine the default direction for service rules. To configure the domain, include the **service-domain** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level:

service-domain (inside | outside);

If you are configuring the interface in a next-hop service-set definition, the **service-domain** setting must match the configuration for the **inside-service-interface** and **outside-service-interface** statements; for more information, see “[Configuring Service Sets to be Applied to Services Interfaces](#)” on page 31.

Related Documentation

- [Configuring Default Timeout Settings for Services Interfaces on page 19](#)
- [Configuring System Logging for Services Interfaces on page 20](#)
- [Examples: Configuring Services Interfaces on page 44](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 608](#)

Configuring System Logging for Service Sets

You specify properties that control how system log messages are generated for the service set. These values override the values configured at the **[edit interfaces interface-name services-options]** hierarchy level.

To configure service-set-specific system logging values, include the **syslog** statement at the **[edit services service-set service-set-name]** hierarchy level:

```
syslog {
  host hostname {
    class class-name
    facility-override facility-name;
    log-prefix prefix-value;
    port port-number
    services severity-level;
    source-address source-address
  }
}
```

Configure the **host** statement with a hostname or an IP address that specifies the system log target server. The hostname **local** directs system log messages to the Routing Engine. For external system log servers, the hostname must be reachable from the same routing instance to which the initial data packet (that triggered session establishment) is delivered. You can specify only one system logging hostname. The **source-address** parameter is supported on the ms, rms, and mams interfaces.

Table 6 on page 47 lists the severity levels that you can specify in configuration statements at the **[edit services service-set service-set-name syslog host hostname]** hierarchy level. The levels from **emergency** through **info** are in order from highest severity (greatest effect on functioning) to lowest.

Table 6: System Log Message Severity Levels

Severity Level	Description
any	Includes all severity levels
emergency	System panic or other condition that causes the router to stop functioning
alert	Conditions that require immediate correction, such as a corrupted system database
critical	Critical conditions, such as hard drive errors
error	Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels
warning	Conditions that warrant monitoring
notice	Conditions that are not errors but might warrant special handling

Table 6: System Log Message Severity Levels (*continued*)

Severity Level	Description
info	Events or non-error conditions of interest

We recommend setting the system logging severity level to **error** during normal operation. To monitor PIC resource usage, set the level to **warning**. To gather information about an intrusion attack when an intrusion detection system error is detected, set the level to **notice** for a specific service set. To debug a configuration or log NAT functionality, set the level to **info**.

For more information about system log messages, see the *Junos OS System Log Messages Reference*.

To select the class of messages to be logged to the specified system log host, include the **class** statement at the `[edit services service-set service-set-name syslog host hostname]` hierarchy level:

```
class class-name;
```

To use one particular facility code for all logging to the specified system log host, include the **facility-override** statement at the `[edit services service-set service-set-name syslog host hostname]` hierarchy level:

```
facility-override facility-name;
```

The supported facilities are: **authorization**, **daemon**, **ftp**, **kernel**, **user**, and **local0** through **local7**.

To specify a text prefix for all logging to this system log host, include the **log-prefix** statement at the `[edit services service-set service-set-name syslog host hostname]` hierarchy level:

```
log-prefix prefix-value;
```

Related Documentation

- [Understanding Service Sets on page 29](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 31](#)
- [Tracing Services PIC Operations on page 48](#)

Tracing Services PIC Operations

Tracing operations track all adaptive services operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

By default, no events are traced. If you include the **traceoptions** statement at the `[edit services adaptive-services-pics]` or `[edit services logging]` hierarchy level, the default tracing behavior is the following:

- Important events are logged in a file called **serviced** located in the `/var/log` directory.

- When the file **serviced** reaches 128 kilobytes (KB), it is renamed **serviced.0**, then **serviced.1**, and so on, until there are three trace files. Then the oldest trace file (**serviced.2**) is overwritten. (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)
- Log files can be accessed only by the user who configures the tracing operation.

You cannot change the directory (**/var/log**) in which trace files are located. However, you can customize the other trace file settings by including the following statements:

```
file filename <files number> <match regular-expression> <size size> <world-readable |
no-world-readable>;
flag {
  all;
  command-queued;
  config;
  handshake;
  init;
  interfaces;
  mib;
  removed-client;
  show;
}
```

You include these statements at the **[edit services adaptive-services-pics traceoptions]** or **[edit services logging traceoptions]** hierarchy level.

These statements are described in the following sections:

- [Configuring the Adaptive Services Log Filename on page 49](#)
- [Configuring the Number and Size of Adaptive Services Log Files on page 49](#)
- [Configuring Access to the Log File on page 50](#)
- [Configuring a Regular Expression for Lines to Be Logged on page 50](#)
- [Configuring the Trace Operations on page 50](#)

Configuring the Adaptive Services Log Filename

By default, the name of the file that records trace output is **serviced**. You can specify a different name by including the **file** statement at the **[edit services adaptive-services-pics traceoptions]** or **[edit services logging traceoptions]** hierarchy level:

```
file filename;
```

Configuring the Number and Size of Adaptive Services Log Files

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed **filename.0**, then **filename.1**, and so on, until there are three trace files. Then the oldest trace file (**filename.2**) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit services adaptive-services-pics traceoptions]** or **[edit services logging traceoptions]** hierarchy level:

```
file <filename> files number size size;
```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (*filename*) reaches 2 MB, *filename* is renamed *filename.0*, and a new file called *filename* is created. When the new *filename* reaches 2 MB, *filename.0* is renamed *filename.1* and *filename* is renamed *filename.0*. This process repeats until there are 20 trace files. Then the oldest file (*filename.19*) is overwritten by the newest file (*filename.0*).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

Configuring Access to the Log File

By default, log files can be accessed only by the user who configures the tracing operation.

To specify that any user can read all log files, include the **file world-readable** statement at the **[edit services adaptive-services-pics traceoptions]** or **[edit services logging traceoptions]** hierarchy level:

```
file <filename> world-readable;
```

To explicitly set the default behavior, include the **file no-world-readable** statement at the **[edit services adaptive-services-pics traceoptions]** or **[edit services logging traceoptions]** hierarchy level:

```
file <filename> no-world-readable;
```

Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including the **match** statement at the **[edit services adaptive-services-pics traceoptions file filename]** or **[edit services logging traceoptions]** hierarchy level and specifying a regular expression (regex) to be matched:

```
file <filename> match regular-expression;
```

Configuring the Trace Operations

By default, if the **traceoptions** configuration is present, only important events are logged. You can configure the trace operations to be logged by including the following statements at the **[edit services adaptive-services-pics traceoptions]** or **[edit services logging traceoptions]** hierarchy level:

```
flag {  
  all;  
  configuration;  
  routing-protocol;  
  routing-socket;  
  snmp;  
}
```

[Table 7 on page 51](#) describes the meaning of the adaptive services tracing flags.

Table 7: Adaptive Services Tracing Flags

Flag	Description	Default Setting
all	Trace all operations.	Off
command-queued	Trace command enqueue events.	Off
config	Log reading of the configuration at the [edit services] hierarchy level.	Off
handshake	Trace handshake events.	Off
init	Trace initialization events.	Off
interfaces	Trace interface events.	Off
mib	Trace GGSN SNMP MIB events.	Off
removed-client	Trace client cleanup events.	Off
show	Trace CLI command servicing.	Off

To display the end of the log, issue the **show log serviced | last** operational mode command:

```
[edit]
user@host# run show log serviced | last
```

Related Documentation

- [Understanding Service Sets on page 29](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 31](#)
- [Configuring System Logging for Service Sets on page 47](#)

PART 3

Translating IP Addresses Using NAT

- [NAT Overview on page 55](#)
- [NAT Configuration Overview on page 65](#)
- [Avoiding IPv4 Exhaustion Using Junos Address Aware Network Addressing and Stateful NAT64 on page 87](#)
- [Hiding Private Networks Using Static Source NAT on page 93](#)
- [Making Private Servers Available Using Static Destination NAT on page 111](#)
- [Allowing Components of a Private Network to Share a Single Address Using NAT on page 117](#)
- [Securing Traffic Using NAT-PT and ALGs on page 141](#)
- [Reducing Traffic and Bandwidth Requirements Using Port Control Protocol on page 165](#)
- [Automatically Assigning Ports Using Port Block Allocation on page 177](#)
- [Connecting Specific Ports and Addresses Using Port Forwarding on page 183](#)
- [Allocating a Few Public Addresses to Many Private Hosts Using Dynamic NAT on page 191](#)
- [Achieving Line-Rate, Low-Latency Translations Using Inline NAT on page 199](#)
- [Monitoring NAT on page 209](#)

CHAPTER 5

NAT Overview

- [Junos Address Aware Network Addressing Overview on page 55](#)
- [Junos OS Carrier-Grade NAT Implementation Overview on page 60](#)
- [Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 61](#)

Junos Address Aware Network Addressing Overview

Junos Address Aware Network Addressing provides Network Address Translation (NAT) functionality for translating IP addresses. It supports a wide range of networking goals, including concealing a set of host addresses on a private network behind a pool of public addresses and providing a security measure to protect the host addresses from direct targeting in network attacks.

Junos Address Aware Network Addressing also provides a tool set to deal with IPv4 exhaustion avoidance, IPv4-IPv6 coexistence, and IPv6 transition technologies. This is particularly important because the Internet Assigned Numbers Authority (IANA) allocated the last large block of IPv4 addresses in early 2011. Service providers, large enterprises, cloud providers, e-tailers, and federal agencies can use Junos Address Aware Network Addressing to pragmatically transition to IPv6 based on business requirements and ensure uninterrupted subscriber and service growth.

- [NAT Concept and Facilities Overview on page 56](#)
- [IPv4-to-IPv4 Basic NAT on page 57](#)
- [Static Destination NAT on page 57](#)
- [Twice NAT on page 57](#)
- [IPv6 NAT on page 58](#)
- [Application-Level Gateway \(ALG\) Support on page 58](#)
- [NAT-PT with DNS ALG on page 58](#)
- [Dynamic NAT on page 59](#)
- [Stateful NAT64 on page 59](#)
- [Dual-Stack Lite on page 59](#)
- [Junos Address Aware Network Addressing Line Card Support on page 60](#)

NAT Concept and Facilities Overview

Junos Address Aware Network Addressing provides carrier-grade NAT (CGN) for IPv4 and IPv6 networks, and facilitates the transit of traffic between different types of networks.

Junos Address Aware Network Addressing supports a diverse set of NAT translation options:

- Static-source translation—Allows you to hide a private network. It features a one-to-one mapping between the original address and the translated address; the mapping is configured statically. For more information, see [“Basic NAT” on page 57](#).
- Dynamic-source translation—Includes two options: dynamic address-only source translation and Network Address Port Translation (NAPT):
 - Dynamic address-only source translation—A NAT address is picked up dynamically from a source NAT pool and the mapping from the original source address to the translated address is maintained as long as there is at least one active flow that uses this mapping. For more information, see [“Dynamic NAT” on page 59](#).
 - NAPT—Both the original source address and the source port are translated. The translated address and port are picked up from the corresponding NAT pool. For more information, see [“NAPT” on page 57](#).
- Static destination translation—Allows you to make selected private servers accessible. It features a one-to-one mapping between the translated address and the destination address; the mapping is configured statically. For more information, see [“Static Destination NAT” on page 57](#).
- Protocol translation—Allows you to assign addresses from a pool on a static or dynamic basis as sessions are initiated across IPv4 or IPv6 boundaries. For more information, see [“Configuring Protocol Translation Between IPv6 and IPv4 Networks - NAT-PT” on page 143](#), [“NAT-PT with DNS ALG” on page 58](#), and [“Stateful NAT64” on page 59](#).
- Encapsulation of IPv4 packets into IPv6 packets using softwires—Enables packets to travel over softwires to a carrier-grade NAT endpoint where they undergo source-NAT processing to hide the original source address. For more information, see [“Tunneling Services for IPv4-to-IPv6 Transition Overview” on page 213](#).

Junos Address Aware Network Addressing supports NAT functionality described in IETF RFCs and Internet drafts, as shown in *“Supported NAT and SIP Standards”* in *Standards Reference*.



NOTE: Not all types of NAT are supported on all interface types. See [“Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card” on page 61](#), which lists features available on supported interfaces.

IPv4-to-IPv4 Basic NAT

Basic Network Address Translation or Basic NAT is a method by which IP addresses are mapped from one group to another, transparent to end users. Network Address Port Translation or NAPT is a method by which many network addresses and their TCP/UDP ports are translated into a single network address and its TCP/UDP ports. Together, these two operations, referred to as traditional NAT, provide a mechanism to connect a realm with private addresses to an external realm with globally unique registered addresses.

Traditional NAT, specified in RFC 3022, *Traditional IP Network Address Translator*, is fully supported by Junos Address Aware Network Addressing. In addition, NAPT is supported for source addresses.

Basic NAT

With Basic NAT, a block of external addresses is set aside for translating addresses of hosts in a private domain as they originate sessions to the external domain. For packets outbound from the private network, Basic NAT translates source IP addresses and related fields such as IP, TCP, UDP, and ICMP header checksums. For inbound packets, Basic NAT translates the destination IP address and the checksums listed above.

NAPT

Use NAPT to enable the components of the private network to share a single external address. NAPT translates the transport identifier (for example, TCP port number, UDP port number, or ICMP query ID) of the private network into a single external address. NAPT can be combined with Basic NAT to use a pool of external addresses in conjunction with port translation.

For packets outbound from the private network, NAPT translates the source IP address, source transport identifier (TCP/UDP port or ICMP query ID), and related fields, such as IP, TCP, UDP, and ICMP header checksums. For inbound packets, NAPT translates the destination IP address, the destination transport identifier, and the IP and transport header checksums.

Static Destination NAT

Use static destination NAT to translate the destination address for external traffic to an address specified in a destination pool. The destination pool contains one address and no port configuration.

For more information about static destination NAT, see RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*.

Twice NAT

In Twice NAT, both the source and destination addresses are subject to translation as packets traverse the NAT router. The source information to be translated can be either address only or address and port. For example, you would use Twice NAT when you are connecting two networks in which all or some addresses in one network overlap with

addresses in another network (whether the network is private or public). In traditional NAT, only one of the addresses is translated.

To configure Twice NAT, you must specify both a destination address and a source address for the match direction, pool or prefix, and translation type.

You can configure application-level gateways (ALGs) for ICMP and traceroute under stateful firewall, NAT, or class-of-service (CoS) rules when Twice NAT is configured in the same service set. These ALGs cannot be applied to flows created by the Packet Gateway Control Protocol (PGCP). Twice NAT does not support other ALGs. By default, the Twice NAT feature can affect IP, TCP, and UDP headers embedded in the payload of ICMP error messages.

Twice NAT, specified in RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*, is fully supported by Junos Address Aware Network Addressing.

IPv6 NAT

IPv6-to-IPv6 NAT (NAT66), defined in Internet draft draft-mrw-behave-nat66-01, *IPv6-to-IPv6 Network Address Translation (NAT66)*, is fully supported by Junos Address Aware Network Addressing.

Application-Level Gateway (ALG) Support

Junos Address Aware Network Addressing supports a number of ALGs. You can use NAT rules to filter incoming traffic based on ALGS. For more information, see [“Network Address Translation Rules Overview” on page 69](#)

NAT-PT with DNS ALG

NAT-PT and Domain Name System (DNS) ALG are used to facilitate communication between IPv6 hosts and IPv4 hosts. Using a pool of IPv4 addresses, NAT-PT assigns addresses from that pool to IPv6 nodes on a dynamic basis as sessions are initiated across IPv4 or IPv6 boundaries. Inbound and outbound sessions must traverse the same NAT-PT router so that it can track those sessions. RFC 2766, *Network Address Translation - Protocol Translation (NAT-PT)*, recommends the use of NAT-PT for translation between IPv6-only nodes and IPv4-only nodes, and *not* for IPv6-to-IPv6 translation between IPv6 nodes or IPv4-to-IPv4 translation between IPv4 nodes.

DNS is a distributed hierarchical naming system for computers, services, or any resource connected to the Internet or a private network. The DNS ALG is an application-specific agent that allows an IPv6 node to communicate with an IPv4 node and vice versa.

When DNS ALG is employed with NAT-PT, the DNS ALG translates IPv6 addresses in DNS queries and responses to the corresponding IPv4 addresses and vice versa. IPv4 name-to-address mappings are held in the DNS with “A” queries. IPv6 name-to-address mappings are held in the DNS with “AAAA” queries.




NOTE: For IPv6 DNS queries, use the `do-not-translate-AAAA-query-to-A-query` statement at the `[edit applications application application-name]` hierarchy level.

Dynamic NAT

Dynamic NAT flow is shown in [Figure 2 on page 59](#).

Figure 2: Dynamic NAT Flow




With dynamic NAT, you can map a private IP address (source) to a public IP address drawing from a pool of registered (public) IP addresses. NAT addresses from the pool are assigned dynamically. Assigning addresses dynamically also allows a few public IP addresses to be used by several private hosts, in contrast with an equal-sized pool required by source static NAT.

For more information about dynamic address translation, see RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*.

Stateful NAT64

Stateful NAT64 flow is shown in [Figure 3 on page 59](#).

Figure 3: Stateful NAT64 Flow



Stateful NAT64 is a mechanism to move to an IPv6 network and at the same time deal with IPv4 address depletion. By allowing IPv6-only clients to contact IPv4 servers using unicast UDP, TCP, or ICMP, several IPv6-only clients can share the same public IPv4 server address. To allow sharing of the IPv4 server address, NAT64 translates incoming IPv6 packets into IPv4 (and vice versa).


When stateful NAT64 is used in conjunction with DNS64, no changes are usually required in the IPv6 client or the IPv4 server. DNS64 is out of scope of this document because it is normally implemented as an enhancement to currently deployed DNS servers.

Stateful NAT64, specified in RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*, is fully supported by Junos Address Aware Network Addressing.

Dual-Stack Lite

Dual-stack lite (DS-Lite) flow is shown in [Figure 4 on page 60](#).

Figure 4: DS-Lite Flow



DS-Lite employs IPv4-over-IPv6 tunnels to cross an IPv6 access network to reach a carrier-grade IPv4-IPv4 NAT. This facilitates the phased introduction of IPv6 on the Internet by providing backward compatibility with IPv4.

Junos Address Aware Network Addressing Line Card Support

Junos Address Aware Network Addressing technologies are available on the following line cards:

- MultiServices Dense Port Concentrator (MS-DPC)
- MS-100, MS-400, and MS-500 MultiServices PICS
- MultiServices Modular Port Concentrator (MS-MPC) and MultiServices Modular Interface Card (MS-MIC)
- Modular Port Concentrator Types 1, 2, and 3 (inline NAT).

Related Documentation

- [Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 61](#)
- [ALGs Available by Default for Junos OS Address Aware NAT on page 141](#)

Junos OS Carrier-Grade NAT Implementation Overview

Junos OS enables you to implement and scale a Carrier-Grade Network Address Translation (CGNAT) solution based on the type of services interfaces used for your implementation:

- MultiServices Denser Port Concentrator (MS-DPC)—The layer 3 services package is used to configure NAT for MS-DPC adaptive services PICs. You must configure the layer-3 services package before implementing NAT on the MS-DPC. This solution provides the NAT functionality described in *Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards*.
- MultiServices Modular Port Concentrator (MS-MPC) and MultiServices Modular Interface Card (MS-MIC)—MS-MPCs and MS-MICs are pre-configured to enable configuration of carrier-grade NAT. This solution provides the NAT functionality described in *Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards*.

- Inline NAT for Type 1, 2, and 3 Modular Port Concentrator (MPC Line Cards)—Inline NAT leverages the services capabilities of TRIO-based MPC line cards, allowing a cost-effective implementation of NAT functionality on the data plane, as described in [“Inline Network Address Translation Overview for MPC Types 1, 2, and 3” on page 199](#).

- Related Documentation**
- [Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 61](#)
 - [Carrier-Grade NAT Implementation: Best Practices on page 76](#)
 - [Example: Configuring Basic NAT44 on page 103](#)

Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card

[Table 8 on page 61](#) summarizes feature differences among the Junos OS carrier-grade NAT implementations.

Table 8: Carrier-Grade NAT—Feature Comparison by Platform

Feature	MS-DPC MS-100 MS-400 MS-500	MS-MPC MS-MIC	MPC Types 1, 2, 3 <i>Inline NAT</i>
Static Source NAT	yes	yes	yes
Dynamic Source NAT - Address Only	yes	yes	no
Dynamic Source NAT - NAPT Port Translation with Secured Port Block Allocation	yes	no	no
Dynamic Source NAT - NAPT Port Translation with Deterministic Port Port Block Allocation	yes	no	no
Static Destination NAT	yes	yes	yes

NOTE: Destination NAT can be implemented indirectly. See [“Inline Network Address Translation Overview for MPC Types 1, 2, and 3” on page 199](#)

Table 8: Carrier-Grade NAT—Feature Comparison by Platform (*continued*)

Feature	MS-DPC MS-100 MS-400 MS-500	MS-MPC MS-MIC	MPC Types 1, 2, 3 <i>Inline NAT</i>
Twice NAT	yes	no	yes <small>NOTE: Twice NAT can be implemented indirectly. See "Inline Network Address Translation Overview for MPC Types 1, 2, and 3" on page 199</small>
NAPT - Preserve Parity and Port	yes	no	no
NAPT - EIM/EIF/APP	yes	yes	no
NAT64	yes	yes	no
NAT64 with APP/EIM/EIF	no	yes	no
NAT64 with ALGs	no	yes	no
<ul style="list-style-type: none"> • FTP • TFTP • SIP • RTSP • PPPT 			
DS-Lite	yes	no	no
6rd	yes	no	no
Overload Pool/Overlap Address Across NAT Pool	yes	no	no
Port Control Protocol	yes	no	no
CGN-PIC	yes	no	no
AMS Support	no	yes	no
Port forwarding	yes	no	no

Table 8: Carrier-Grade NAT—Feature Comparison by Platform (*continued*)

Feature	MS-DPC MS-100 MS-400 MS-500	MS-MPC MS-MIC	MPC Types 1, 2, 3 <i>Inline NAT</i>
No translation	yes	yes	yes
NOTE: The no-translation statement is not supported for ms- interfaces on MS-MICs and MS-MPCs.			
No translation	yes	no	yes
No translation	yes	no	yes

Table 9 on page 63 summarizes availability of translation types by type of line card.

Table 9: Carrier-Grade NAT Translation Types

Translation Type	MS-DPC MS-100 MS-400 MS-500	MS-MPC MS-MIC	MPC Types 1, 2, 3 <i>Inline NAT</i>
basic-nat44	yes	yes	yes
basic-nat66	yes	no	no
basic-nat-pt	yes	no	no
deterministic-napt44	yes	no	no
dnat-44	yes	yes	no
dynamic-nat44	yes	yes	no
napt-44	yes	yes	no
napt-66	yes	no	no
napt-pt	yes	no	no
stateful-nat64	yes	yes	no

Table 9: Carrier-Grade NAT Translation Types (*continued*)

Translation Type	MS-DPC		
	MS-100		
	MS-400	MS-MPC	MPC Types 1, 2, 3
	MS-500	MS-MIC	<i>Inline NAT</i>
twice-basic-nat-44	yes	no	yes
twice-dynamic-nat-44	yes	no	no
twice-dynamic-napt-44	yes	no	no

Related Documentation

- [Junos OS Carrier-Grade NAT Implementation Overview on page 60](#)

CHAPTER 6

NAT Configuration Overview

- [Network Address Translation Configuration Overview on page 65](#)
- [Configuring Source and Destination Addresses Network Address Translation Overview on page 66](#)
- [Configuring Pools of Addresses and Ports for Network Address Translation Overview on page 67](#)
- [Network Address Translation Rules Overview on page 69](#)
- [Configuring Service Sets for Network Address Translation on page 75](#)
- [Carrier-Grade NAT Implementation: Best Practices on page 76](#)

Network Address Translation Configuration Overview

To configure network address translation (NAT), complete the following high-level steps:

1. Configure the source and destination addresses. For more information, see [“Configuring Source and Destination Addresses Network Address Translation Overview” on page 66](#).
2. Define the addresses or prefixes, address ranges, and ports used for NAT. For more information, see [“Configuring Pools of Addresses and Ports for Network Address Translation Overview” on page 67](#)
3. If applicable, configure the address pools for network address port translation (NAPT). For more information, see [“Configuring Address Pools for Network Address Port Translation \(NAPT\) Overview” on page 117](#).
4. Configure the NAT rules. Within the rules, include match directions, match conditions, actions, and translation types. For more information, see [“Network Address Translation Rules Overview” on page 69](#).
5. Configure service sets for NAT processing. Within each service set, define the interfaces for handling inbound and outbound traffic and a NAT rule or ruleset. For more information, see [“Configuring Service Sets for Network Address Translation” on page 75](#).

Related Documentation

- [Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards](#)
- [Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 61](#)

Configuring Source and Destination Addresses Network Address Translation Overview

You must configure a specific address, a prefix, or the address-range boundaries:

- The following addresses, while valid in **inet.0**, cannot be used for NAT translation:
 - **0.0.0.0/32**
 - **127.0.0.0/8** (loopback)
 - **128.0.0.0/16** (martian)
 - **191.255.0.0/16** (martian)
 - **192.0.0.0/24** (martian)
 - **223.255.255.0/24** (martian)
 - **224.0.0.0/4** (multicast)
 - **240.0.0.0/4** (reserved)
 - **255.255.255.255** (broadcast)

The addresses that are specified as valid in the **inet.0** routing table and not supported for NAT translation are **orlonger** match filter types. You cannot specify any regions within such address prefixes in a NAT pool.

- On MX Series routers with MS-MPCs and MS-MICs, if you configure a NAT address pool with a prefix length that is equal to or greater than /16, the PIC does not contain sufficient memory to provision the configured pool. Also, memory utilization problems might occur if you attempt to configure many pools whose combined total IP addresses exceed /16. In such circumstances, a system logging message is generated stating that the NAT pool name is failed to be created and that the service set is not activated. On MS-MPCs and MS-MICs, you must not configure NAT pools with prefix lengths greater than or equal to /16.
- You can specify one or more IPv4 address prefixes in the **pool** statement and in the **from** clause of the NAT rule term. This enables you to configure source translation from a private subnet to a public subnet without defining a rule term for each address in the subnet. Destination translation cannot be configured by this method. For more information, see *Examples: Configuring NAT Rules*.
- When you configure static source NAT, the **address** prefix size you configure at the **[edit services nat pool *pool-name*]** hierarchy level must be larger than the **source-address** prefix range configured at the **[edit services nat rule *rule-name* term *term-name* from]** hierarchy level. The **source-address** prefix range must also map to a single subnet or range of IPv4 or IPv6 addresses in the **pool** statement. Any pool addresses that are not used by the **source-address** prefix range are left unused. Pools cannot be shared.



NOTE: When you include a NAT configuration that changes IP addresses, it might affect forwarding path features elsewhere in your router configuration, such as source class usage (SCU), destination class usage (DCU), filter-based forwarding, or other features that target specific IP addresses or prefixes.

NAT configuration might also affect routing protocol operation, because the protocol peering, neighbor, and interface addresses can be altered when routing protocols packets transit the Adaptive Services (AS) or Multiservices PIC.

Related Documentation

- [Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards](#)

Configuring Pools of Addresses and Ports for Network Address Translation Overview

- [Configuring NAT Pools on page 67](#)
- [Preserve Range and Preserve Parity on page 68](#)
- [Specifying Destination and Source Prefixes without Configuring a Pool on page 68](#)

Configuring NAT Pools

You can use the **pool** statement to define the addresses (or prefixes), address ranges, and ports used for Network Address Translation (NAT). To configure the information, include the **pool** statement at the **[edit services nat]** hierarchy level:

```
[edit services nat]
pool nat-pool-name {
  address ip-prefix</prefix-length>;
  address-range low minimum-value high maximum-value;
  port {
    automatic (sequential | random-allocation);
    range low minimum-value high maximum-value random-allocation;
    preserve-parity;
    preserve-range {
    }
  }
}
```

To configure pools for traditional NAT, specify either a destination pool or a source pool.

With static source NAT and dynamic source NAT, you can specify multiple IPv4 addresses (or prefixes) and IPv4 address ranges. Up to 32 prefixes or address ranges (or a combination) can be supported within a single pool.

With static destination NAT, you can also specify multiple address prefixes and address ranges in a single term. Multiple destination NAT terms can share a destination NAT pool. However, the netmask or range for the **from** address must be smaller than or equal to the netmask or range for the destination pool address. If you define the pool to be larger than required, some addresses will not be used. For example, if you define the pool size as 100 addresses and the rule specifies only 80 addresses, the last 20 addresses in the pool are not used.

For constraints on specific translation types, see [“Network Address Translation Rules Overview” on page 69](#).

With source static NAT, the prefixes and address ranges cannot overlap between separate pools.

In an address range, the **low** value must be a lower number than the **high** value. When multiple address ranges and prefixes are configured, the prefixes are depleted first, followed by the address ranges.

When you specify a port for dynamic source NAT, address ranges are limited to a maximum of 65,000 addresses, for a total of (65,000 x 65,535) or 4,259,775,000 flows. A dynamic NAT pool with no address port translation supports up to 65,535 addresses. There is no limit on the pool size for static source NAT.

Preserve Range and Preserve Parity

You can configure your carrier-grade NAT (CGN) to preserve the range or parity of the packet source port when it allocates a source port for an outbound connection. You can configure the preserve parity and preserve range options under the NAT pool definition by including the **preserve-range** and **preserve-parity** configuration statements at the **[edit services nat pool poolname port]** hierarchy level.

- Preserve range—RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*, defines two ranges: 0 through 1023, and 1024 through 65,535. When the **preserve-range** knob is configured and the incoming port falls into one of these ranges, CGN allocates a port from that range only. However, if there is no available port in the range, the port allocation request fails and that session is not created. The failure is reflected on counters and system logging, but no Internet Control Message Protocol (ICMP) message is generated. If this knob is not configured, allocation is based on the configured port range without regard to the port range that contains the incoming port. The exception is some application-level gateways (ALGs), such as hello, that have special zones.
- Preserve parity—When the **preserve-parity** knob is configured, CGN allocates a port with the same even or odd parity as the incoming port. If the incoming port number is odd or even, the outgoing port number should correspondingly be odd or even. If a port number of the desired parity is not available, the port allocation request fails, the session is not created, and the packet is dropped.

Specifying Destination and Source Prefixes without Configuring a Pool

You can directly specify the destination or source prefix used in NAT without configuring a pool.

To configure the information, include the **rule** statement at the **[edit services nat]** hierarchy level:

```
[edit services nat]
rule rule-name {
  term term-name {
    then {
      translated {
```

```

        destination-prefix prefix;
    }
}
}
}

```

Network Address Translation Rules Overview

To configure a NAT rule, include the **rule** *rule-name* statement at the **[edit services nat]** hierarchy level:

```

[edit services nat]
rule rule-name {
  match-direction (input | output);
  term term-name {
    from {
      application-sets set-name;
      applications [ application-names ];
      destination-address (address | any-unicast) <except>;
      destination-address-range low minimum-value high maximum-value <except>;
      destination-prefix-list list-name <except>;
      source-address (address | any-unicast) <except>;
      source-address-range low minimum-value high maximum-value <except>;
      source-prefix-list list-name <except>;
    }
    then {
      no-translation;
      translated {
        address-pooling paired;
        destination-pool nat-pool-name;
        destination-prefix destination-prefix;
        dns-alg-pool dns-alg-pool;
        dns-alg-prefix dns-alg-prefix;
        filtering-type endpoint-independent;
        mapping-type endpoint-independent;
        overload-pool overload-pool-name;
        overload-prefix overload-prefix;
        source-pool nat-pool-name;
        source-prefix source-prefix;
        translation-type {
          (basic-nat-pt | basic-nat44 | basic-nat66 | dnat-44 | dynamic-nat44 | napt-44 |
           napt-66 | napt-pt | stateful-nat64 | twice-basic-nat-44 | twice-dynamic-nat-44
           | twice-napt-44);
        }
      }
    }
    syslog;
  }
}
}

```

Each rule must include a **match-direction** statement that specifies the direction in which the match is applied.

In addition, each NAT rule consists of a set of terms, similar to a firewall filter. A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded.
- **then** statement—Specifies the actions and action modifiers to be performed by the router software.

The following sections explain how the components of NAT rules:

- [Configuring Match Direction for NAT Rules on page 70](#)
- [Configuring Match Conditions in NAT Rules on page 70](#)
- [Configuring Actions in NAT Rules on page 71](#)
- [Configuring Translation Types on page 73](#)

Configuring Match Direction for NAT Rules

Each rule must include a **match-direction** statement that specifies the direction in which the match is applied. To configure where the match is applied, include the **match-direction** statement at the **[edit services nat rule *rule-name*]** hierarchy level:

```
[edit services nat rule rule-name]  
match-direction (input | output);
```

The match direction is used with respect to the traffic flow through the Multiservices DPC and Multiservices PICs. When a packet is sent to the PIC, direction information is carried along with it. The packet direction is determined based on the following criteria:

- With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.
- With a next-hop service set, packet direction is determined by the interface used to route the packet to the Multiservices DPC or Multiservices PIC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the PIC or DPC, the packet direction is output. For more information about inside and outside interfaces, see [“Configuring Service Sets to be Applied to Services Interfaces” on page 31](#).
- On the Multiservices DPC and Multiservices PIC, a flow lookup is performed. If no flow is found, rule processing is performed. All rules in the service set are considered. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that matches the packet direction are considered.

Configuring Match Conditions in NAT Rules

To configure NAT match conditions, include the **from** statement at the **[edit services nat rule *rule-name* term *term-name*]** hierarchy level:

```
[edit services nat rule rule-name term term-name]  
from {  
  application-sets set-name;  
  applications [ application-names ];  
  destination-address (address | any-unicast) <except>;  
  destination-address-range low minimum-value high maximum-value <except>;  
  destination-prefix-list list-name <except>;  
  source-address (address | any-unicast) <except>;
```



```

source-address-range low minimum-value high maximum-value <except>;
source-prefix-list list-name <except>;
}

```

To configure traditional NAT, you can use the destination address, a range of destination addresses, the source address, or a range of source addresses as a match condition, in the same way that you would configure a firewall filter; for more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Alternatively, you can specify a list of source or destination prefixes by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level and then including either the **destination-prefix-list** or **source-prefix-list** statement in the NAT rule. For an example, see “[Examples: Configuring Stateful Firewall Rules](#)” on page 335.

If the **translation-type** statement in the **then** statement of the nat rule is set to **stateful-nat-64**, the range specified by the **destination-address-range** or the **destination-prefix-list** in the **from** statement must be within the range specified by the **destination-prefix** statement in the **then** statement.

If at least one NAT term within a NAT rule has the address pooling paired (APP) functionality enabled (by including the **address-pooling** statement at the **[edit services nat rule rule-name term term-name then translated]** hierarchy level, all the other terms in the NAT rule that use the same NAT address pool as the address pool for the term with APP enabled must have APP enabled. Otherwise, if you add a NAT rule term without enabling APP to a rule that contains other terms with APP enabled, all the terms with APP enabled in a NAT rule drop traffic flows that match the specified criteria in the NAT rule.

For MX Series routers with MS-MICs and MS-MPCs, although the address pooling paired (APP) functionality is enabled within a NAT rule (by including the **address-pooling** statement at the **[edit services nat rule rule-name term term-name then translated]** hierarchy level), it is a characteristic of a NAT pool. Such a NAT pool for which APP is enabled cannot be shared with NAT rules that do not have APP configured.

Configuring Actions in NAT Rules

To configure NAT actions, include the **then** statement at the **[edit services nat rule rule-name term term-name]** hierarchy level:

```

[edit services nat]
rule rule-name {
  term term-name {
    from {
      destination-address-range low minimum-value high maximum-value <except>;
      destination-prefix-list list-name <except>;
    }
    then {
      destination-prefix destination-prefix;
    }
  }
}

[edit services nat rule rule-name term term-name]
then {
  no-translation;
  syslog;
}

```

```
translated {  
  destination-pool nat-pool-name;  
  destination-prefix destination-prefix;  
  source-pool nat-pool-name;  
  source-prefix source-prefix;  
  translation-type (basic-nat-pt | basic-nat44 | basic-nat66 | dnat-44 | dynamic-nat44  
    | napt-44 | napt-66 | napt-pt | stateful-nat64 | twice-basic-nat-44 |  
    twice-dynamic-nat-44 | twice-napt-44);  
}  
}
```

- The **no-translation** statement allows you to specify addresses that you want excluded from NAT.
- The **system log** statement enables you to record an alert in the system logging facility.
- The **destination-pool**, **destination-prefix**, **source-pool**, and **source-prefix** statements specify addressing information that you define by including the **pool** statement at the **[edit services nat]** hierarchy level; for more information, see [“Configuring Pools of Addresses and Ports for Network Address Translation Overview” on page 67](#).
- The **translation-type** statement specifies the type of NAT used for source or destination traffic. The options are **basic-nat-pt**, **basic-nat44**, **basic-nat66**, **dnat-44**, **dynamic-nat44**, **napt-44**, **napt-66**, **napt-pt**, **stateful-nat64**, **twice-basic-nat-44**, **twice-dynamic-nat-44**, and **twice-napt-44**.



NOTE: In Junos OS Release 13.2 and earlier, the following restriction was not enforced by the CLI: if the **translation-type** statement in the then statement of a NAT rule was set to **stateful-nat-64**, the range specified by the **destination-address-range** or the **destination-prefix-list** in the from statement needed to be within the range specified by the **destination-prefix** statement in the then statement. Starting in Junos OS Release 13.3R1, this restriction is enforced.

Configuring Translation Types

The implementation details of the nine options of the **translation-type** statement are as follows:

- **basic-nat44**—This option implements the static translation of source IP addresses without port mapping. You must configure the **from source-address** statement in the match condition for the rule. The size of the address range specified in the statement must be the same as or smaller than the source pool. You must specify either a source pool or a destination prefix. The referenced pool can contain multiple addresses but you cannot specify ports for translation.



NOTE: In an interface service set, all packets destined for the source address specified in the match condition are automatically routed to the services PIC, even if no service set is associated with the interface.



NOTE: Prior to Junos OS Release 11.4R3, you could only use a source NAT pool in a single service set. As of Junos OS Release 11.4R3 and subsequent releases, you can reuse a source NAT pool in multiple service sets.

- **basic-nat66**—This option implements the static translation of source IP addresses without port mapping in IPv6 networks. The configuration is similar to the **basic-nat44** implementation, but with IPv6 addresses.
- **basic-nat-pt**—This option implements translation of addresses of IPv6 hosts, as they originate sessions to the IPv4 hosts in an external domain and vice versa. This option is always implemented with DNS ALG. You must define the source and destination pools of IPv4 addresses. You must configure one rule and define two terms. Configure the IPv6 addresses in the **from** statement in both **term** statements. In the **then** statement of the first term within the rule, reference both the source and destination pools and configure **dns-alg-prefix**. Configure the source prefix in the **then** statement of the second term within the same rule.
- **dnat-44**—This option implements static translation of destination IP addresses without port mapping. The size of the pool address space must be greater than or equal to the destination address space. You must specify a name for the **destination pool** statement. The referenced pool can contain multiple addresses, ranges, or prefixes, as long as the number of NAT addresses in the pool is larger than the number of destination addresses in the **from** statement. You must include exactly one **destination-address** value at the **[edit services nat rule rule-name term term-name from]** hierarchy level; if it is a prefix, the size must be less than or equal to the pool prefix size. Any addresses in the pool that are not matched in the yvalue remain unused, because a pool cannot be shared among multiple terms or rules.
- **dynamic-nat44**—This option implements dynamic translation of source IP addresses without port mapping. You must specify a **source-pool**. The referenced pool must include an **address** configuration (for address-only translation).

The **dynamic-nat44** address-only option supports translating up to 16,777,216 addresses to a smaller size pool. The requests from the source address range are assigned to the addresses in the pool until the pool is used up, and any additional requests are rejected. A NAT address assigned to a host is used for all concurrent sessions from that host. The address is released to the pool only after all the sessions for that host expire. This feature enables the router to share a few public IP addresses between several private hosts. Because all the private hosts might not simultaneously create sessions, they can share a few public IP addresses.

- **napt-44**—This option implements dynamic translation of source IP addresses with port mapping. You must specify a name for the **source-pool** statement. The referenced pool must include a **port** configuration. If the port is configured as automatic or a port range is specified, then it implies that Network Address Port Translation (NAPT) is used.
- **napt-66**—This option implements dynamic address translation of source IP addresses with port mapping for IPv6 addresses. The configuration is similar to the **napt-44** implementation, but with IPv6 addresses.
- **napt-pt**—This option implements dynamic address and port translation for source and static translation of destination IP address. You must specify a name for the **source-pool** statement. The referenced pool must include a port configuration (for NAPT). Additionally, you must configure two rules, one for the DNS traffic and the other for the rest of the traffic. The rule meant for the DNS traffic should be DNS ALG enabled and the **dns-alg-prefix** statement should be configured. Moreover, the prefix configured in the **dns-alg-prefix** statement must be used in the second rule to translate the destination IPv6 addresses to IPv4 addresses.
- **stateful-nat64**—This option implements dynamic address and port translation for source IP addresses and prefix removal translation for destination IP addresses. You must specify the IPv4 addresses used for translation at the **[edit services nat pool]** hierarchy level. This pool must be referenced in the rule that translates the IPv6 addresses to IPv4.
- **twice-basic-nat-44**—This option implements static source and static destination translation for IPv4 addresses, thus combining **basic-nat44** for source and **dnat-44** for destination addresses.
- **twice-dynamic-nat-44**—This option implements source dynamic and destination static translation for IPv4 addresses, combining **dynamic-nat44** for source and **dnat-44** for destination addresses.
- **twice-napt-44**—This option implements source NAPT and destination static translation for IPv4 addresses, combining **napt-44** for source and **dnat-44** for destination addresses.



NOTE: When configuring NAT, if any traffic is destined for the following addresses and does not match a NAT flow or NAT rule, the traffic is dropped:

- Addresses specified in the `from destination-address` statement when you are using destination translation
- Addresses specified in the source NAT pool when you are using source translation

For more information on NAT methods, see RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*.

Configuring Service Sets for Network Address Translation

When configuring a service set for NAT processing, make sure you have defined:

- Service interface(s) for handling inbound and outbound traffic



NOTE: Prior to Junos OS Release 11.4R3, you could only use a source NAT pool in a single service set. As of Junos OS Release 11.4R3 and subsequent releases, you can reuse a source or destination NAT pool in multiple service sets, provided that the service interfaces associated with the service sets are in different virtual routing and forwarding (VRF) instances.

- For interface style service sets, when a NAT pool is reused in multiple service sets, the service interfaces used in the `interface-service service-interface` option of each service set must be in different VRFs.
- For next-hop style service sets, when a NAT pool is reused in multiple service sets, the service interfaces used in the `outside-interface` option of each service set must be in different VRFs.

Not adhering to these service interface restrictions will cause multiple routes to be installed in the same VRF for the same NAT addresses, causing reverse traffic to be processed incorrectly.

To enable sharing of source NAT pools, include the `allow-overlapping-nat-pools` statement at the `[edit services nat]` hierarchy level.

- A NAT rule or ruleset



NOTE: To configure an MX-DPC interface to be used exclusively for carrier-grade NAT (CGN) or related services (intrusion detection, stateful firewall, and software), include the `cg-pic` statement at the `[edit interfaces interface-name services-options]` hierarchy level.

To configure a NAT service set:

1. At the **[edit services]** hierarchy level, define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

Or

```
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface interface-name
                        outside-service-interface interface-name
```



NOTE: If you have a Trio-based line card (MPC/MIC), you can use an inline-services interface that was configured on that card, as shown in this example:

```
[edit]
user@host# set interfaces si-0/0/0
[edit services service-set s1]
user@host# set interface-service service-interface si-0/0/0
```

For more information on interface service and next-hop service, see [“Configuring Service Sets to be Applied to Services Interfaces” on page 31.](#)

3. Configure a reference to the NAT rules or ruleset to be used with the service set.

```
[edit services service-set service-set-name]
user@host# set nat-rules rule-or-ruleset-name
```

4. (Optional) For NAT64, specify that the don't fragment (DF) bit for IPv4 packet headers is cleared when packet length is less than 1280 bytes.

```
[edit services service-set service-set-name]
user@host# set nat-options stateful-nat64 clear-dont-fragment-bit
```

Related Documentation

- [Configuring Service Sets to be Applied to Services Interfaces on page 31](#)

Carrier-Grade NAT Implementation: Best Practices

The following topics present the best practices for carrier-grade NAT implementation on MS-DPCs using the Layer 3 services package:

- [Use APP and Round-Robin Address-Allocation on page 77](#)
- [Do Not Use EIM with SIP on page 77](#)
- [Do Not Use EIM with HTTP, DNS, or When Not Needed on page 78](#)
- [Define PBA Blocks Based on User Profiles on page 79](#)

- [Do Not Change the PBA Configuration on Running Systems on page 79](#)
- [Do Not Allocate Excessively Large NAT Pools on page 80](#)
- [Configure the System Log for PBA Only When Needed on page 81](#)
- [Use Redundant Service PIC \(RSP\) Interfaces for Failover on page 83](#)
- [Contain the Effects of Missing IP Fragments on page 84](#)
- [Do Not Use Configurations Prone to Routing Loops on page 84](#)

Use APP and Round-Robin Address-Allocation

Scenario:

- Address-pooling paired (APP) allows a private IP address to be mapped to the same public IP address from a NAT pool for all its sessions. The binding between private IP and public IP is triggered by the first packet seen from such private host.
- By default, an MS-DPC or MS-PIC allocates ports from a NAT pool in a sequential fashion from each consecutive IP address available in the pool.
- Sequential allocation, together with APP, can result in mapping multiple private hosts to the same public IP address, resulting in fast port exhaustion for the interested public IP address while other ports are still available from the remaining of NAT pool.



BEST PRACTICE: Configure round-robin address allocation for the NAT pool used by traffic served with APP. Round-robin allocation allocates ports from different IP addresses.

The following snippet provides an example of round-robin address allocation.

```
user@router# show services nat pool natpool-1
address-range low 9.9.9.1 high 9.9.9.10;
port {
    automatic;
}
address-allocation round-robin;
mapping-timeout 120;
```

Do Not Use EIM with SIP

Scenario:

- Session Initiation Protocol (SIP) traffic requires an Application Level Gateway (ALG) to allow SIP servers and clients on the public side of the CGNAT to communicate with the SIP hosts on the private side.
- The SIP ALG opens the pinholes in the CGNAT router to permit the forwarding of outbound traffic based on any supported SIP feature.
- Endpoint-independent mapping (EIM) is not needed by SIP to function, nor by the SIP ALG to create the flows for forwarding the SIP traffic



BEST PRACTICE: Do *not* configure EIM together with the SIP ALG; doing so adds processing overhead with no benefit.

```
user@router# show services nat rule natrule-1
match-direction input;
term 1 {
  from {
    applications junos-sip;
  }
  then {
    translated {
      source-pool natpool-3;
      translation-type {
        napt-44;
      }
    }
    address-pooling paired;
  }
}
```

Do Not Use EIM with HTTP, DNS, or When Not Needed

Scenario:

- Most Internet traffic uses HTTP, and there is no browser on any OS that reuses the same source port for sending traffic to different destinations. EIM provides no benefit for HTTP traffic.
- Because none of the junos-algs require EIM to work, avoid using EIM with the ALGs.
- EIM allocates memory for each mapping; this is in addition to the memory used for flow allocation. This reduces the maximum number of flows that can be established through the services PIC, and causes processing overhead for the creation and deletion of flows and mappings.



BEST PRACTICE:

- Don't enable EIM for applications that are defined ALGs or are known not to use Session Traversal Utilities for NAT (STUN) servers to discover the presence of a NAT router.
- Enable EIM for applications that do reuse the source ports and rely on a CGNAT device to maintain the same address:port mapping for all traffic sent to different destinations, such as on-line gaming applications like Xbox and PS3, or applications that use unilateral self-address fixing methods (UNSAF). see (*IETF RFC 3424 IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation*).

Define PBA Blocks Based on User Profiles

Scenario:

- When a user connects to a website that requires the establishment of a significant number of sockets for a single HTML page, a corresponding number of new ports must be allocated. Port blocks should be large enough to prevent continual allocation of new blocks.
- If the number of concurrent sessions exceeds the number of ports available in the active port block, the other allocated port-blocks will be scanned for available ports to use or a new block will be allocated from the free block pool.
- The process of continually scanning the allocated port-blocks and/or allocating additional blocks from the free block pool could result in experienced latency for setting up new sessions and delay loading of web pages.
- Having a user continuously allocating or de-allocating from different PBA blocks impacts performance.



BEST PRACTICE: Define PBA blocks with a size that is a power of 2 or 4 related to the average number of sessions a user is expected to have active. For example, if a user is expected to have an average of approximately 200 to 250 sessions active, configuring the PBA block size to 512 or 1024 will provide a liberal allocation.

```
user@router# show services nat pool natpool-1
address-range low 9.9.9.1 high 9.9.9.10;
  port {
    automatic;
    secure-port-block-allocation {
      block-size 1024;
      max-blocks-per-user 8; /* Max 2048, default 8 */
      active-block-timeout 300;
    }
  }
mapping-timeout 300;
```

Do Not Change the PBA Configuration on Running Systems

Scenario:

- PBA settings in NAT pools are mapped to memory at the time of the Service PIC boot up and cannot be changed while processing traffic.
- Do not change the following settings:
 - Update any NAT pool PBA configuration.
 - Change a PBA NAT pool to a non-PBA NAT pool.
 - Change a non-PBA NAT pool to a PBA NAT pool.

Any of these changes result in the logging of the following message:

PBA_CATASTROPIC_CHANGE: The recent PBA configuration changes will reflect in the Service-PIC only after deactivate and activate of the service-set again



NOTE: Starting with Junos OS Release 14.1, when you deactivate a service-set that contains address pooling paired (APP) or endpoint independent mapping (EIM) mapping for that service-set, messages are displayed on the PIC console and the mappings are cleared for that service-set. These messages are triggered when the deletion of a service-set commences and again generated when the deletion of the service-set is completed. The following sample messages are displayed when deletion starts and ends:

- Nov 15 08:33:13.974 LOG: Critical] SVC-SET ss1 (iid 5) deactivate/delete: NAT Mappings and flows deletion initiated
- Nov 15 08:33:14.674 LOG: Critical] SVC-SET ss1 (iid 5) deactivate/delete: NAT Mappings and flows deletion completed

In a scaled environment that contains a large number of APP or EIM mappings in a service set, a heavy volume of messages is generated and this process takes some amount of time. We recommend that you wait until the console messages indicating the completion of deletion of the service set are completed before you reactivate the service-set again.



BEST PRACTICE: When changing PBA configurations, restart the services PIC if possible. Minimally, you must deactivate and reactivate the affected service set.

Do Not Allocate Excessively Large NAT Pools

Scenario:

- The maximum number of flows supported by the MS-DPC and each PIC on an MS-DPC is 8 million.
- Assuming that the 8 million flow maximum consists of 4 million sessions (1 reverse flow for each forward flow), these sessions would require a maximum of 4 million ports that are available from 64 IP addresses within the 1024 to 65,535 ports range (64K ports per IP address).
- Do not configure ports to support more than 8 million flows; they will never be needed.
- This scenario assumes that APP, EIM, and EIF are not enabled. When they *are* enabled, the total number of flows is lower, which means that you should configure the number of IP addresses in the NAT pool based on the maximum supported flows.



BEST PRACTICE: Do not configure NAT pools with more than 64 addresses (that is, a /26 network) and round-robin configured and 64K ports from each address.

On MS-MICs, do not configure NAT pools with more than 128 addresses (that is, a /25 network) and round-robin configured and 64K ports from each address. On MS-MICs, a maximum of up to 7 million sessions are supported. Assuming these 7 million sessions, such sessions would require maximum 7 million ports that are available from 128 IP addresses within the 1024-65535 ports range (64K ports per IP address).

On MS-MPCs, do not configure NAT pools with more than 256 addresses (that is, a /24 network) and round-robin configured and 64K ports from each address. On MS-MPCs, a maximum of up to 15 million sessions are supported. Assuming these 15 million sessions would require maximum 15 million ports that are available from 256 IP addresses within the 1024-65535 ports range (64K ports per IP address).

Configure the System Log for PBA Only When Needed

Scenario:

- Session logging can negatively affect performance depending on the frequency of creation and deletion of flows.
- PBA is meant to reduce the need for logging.
- Deterministic NAT is designed to eliminate the need for logging.
- All system log messages created by the services PIC constitutes traffic that will be sent to the Packet Forwarding Engine, competing with user traffic to reach the external destination.



BEST PRACTICE:

- Use logging to the system log at the service-set level rather than at the services PIC interface level when possible.
- Do not enable logging for redundant information. When using PBA, you don't need to configure logs per session because knowing the PBA block and the block size enables you to derive the ports allocated to each user. In this case, a log that reports all sessions created by that user with ports belonging to a block is redundant. If you have configured deterministic NAT (DetNat) a log is completely unnecessary because all information on port allocation can be deduced mathematically.
- Rate-limit the number of logs generated from an sp- interface. When not set, the default limits apply: 10K for the local host system log server (RE) and 200K for the external system log server.

```
user@router# show interfaces sp-1/1/0 services-options
```

```
system log {
  host 1.2.3.4 {
    services info;
  }
  message-rate-limit 1000;
}
```

- Always system log to an external server to avoid loading the Routing Engine and specify system log class to restrict logging.
 - If you do not specify system log class, all log messages are allowed (subject to priority check and rate limiting).
 - When you specify system log class, only messages meeting the class criteria are retained.
 - Use the `show services service-sets statistics system log detail` command to check what is being dropped by unconfigured classes.

```
user@router# show services service-set S-SET-1 system log
host 1.2.3.4 {
  services info;
  class {
    session-logs open close;
    packet-logs;
    stateful-firewall-logs;
    alg-logs;
    nat-logs;
    ids-logs;
  }
}
```



BEST PRACTICE: System log generation can be *rule-based* or *event-based*.

- Use rule-based system logging with care; it generates a log for every packet that enters the rule term, since rule-based logging is not subject to class or priority filtering.
- System log messages can be dropped only as a result of message rate limiting. Make sure you have set a realistic rate-limit that is unlikely to be exceeded.
- Use rule-based logging only for discarded traffic (a relatively small percentage of the traffic) or for troubleshooting. Since rule-based logging applies to all traffic that enters the PIC and creates a flow, logging can be excessive, resulting in reaching the configured induce rate limit with a consequent loss of needed logs.

```
cli# show services stateful-firewall
rule rule-sfw-accept {
  match-direction input-output;
  term term-sfw-accept {
    then {
      accept;
      system log;
    }
  }
}
```

```

}
rule rule-sfw-reject {
  match-direction input-output;
  term term-sfw-reject {
    then {
      reject;
      system log;
    }
  }
}

```

**BEST PRACTICE:**

All rule match logs are enabled by their respective rules:

- ASP_COS_RULE_MATCH (class-of-service rules)
- ASP_COS_RULE_MATCH (class-of-service rules)
- ASP_IDS_RULE_MATCH (ids rules)
- ASP_NAT_RULE_MATCH (nat rule)
- ASP_SFW_RULE_ACCEPT (stateful firewall rules)
- ASP_SFW_RULE_DISCARD
- ASP_SFW_RULE_REJECT

Use Redundant Service PIC (RSP) Interfaces for Failover

**BEST PRACTICE:**

- The usage of Redundant Service PIC (RSP) interfaces, allows the active services PIC to perform an immediated switchover to the secondary services PIC in case of major issues that require a services PIC reboot.
- This results in a minimal service impact for user traffic.
- There are two modes for redunancy: warm-standby (default) and hot-standby. Hot-standby provides 1:1 redundancy, while warm-standby provides 1:N redundancy. With both modes , there is no impact on the UDP forwarding.
- When the secondary services PIC is shared among multiple RSPs, only warm-standby is possible and the impact to traffic is limited to the time to load the appropriate configuration on the secondary PIC.

```

user@router# show interfaces rsp0
redundancy-options {
  primary sp-0/1/0;
  secondary sp-1/1/0;
  hot-standby;
}

```

Contain the Effects of Missing IP Fragments

Scenario:

- IP fragments are buffered as they arrive to facilitate the integrity check of the completely reassembled packet before being serviced by the services PIC.
- Missing fragments cause received fragments to be held until the internal buffer is full and are flushed out. This causes CPU usage overhead and reduced traffic forwarding.



BEST PRACTICE: Configure the `fragment-limit`, the maximum number of fragments for a packet, and `reassembly-timeout`, the maximum wait for a missing fragment, after which all other fragments for the same packet are flushed out.

```
user@router# show interfaces sp-0/0/0
services-options {
    open-timeout 5;
    close-timeout 5;
    inactivity-timeout 30;
    tcp-tickles 4;
    fragment-limit 10;
    reassembly-timeout 3;
    cgn-pic;
}
```

Do Not Use Configurations Prone to Routing Loops

Scenario:

- Sudden and persistent high CPU usage is most likely an indication of packet looping between the Packet Forwarding Engine and the services PIC. Depending on whether the configuration uses interface-style or next-hop-style service sets, different network flaps can lead to routing loops.



BEST PRACTICE:

Ensure that only the intended traffic is allowed to reach the services PIC and is serviced based on service set rule.

- Configure a firewall filter that accepts only the traffic meant to go to the services PIC on the output direction of the `sp-` interface. That is, accept only traffic identified in the NAT rule from option as received from the `source-address` that identifies the customer private network; discard and log all the rest.
- Allow only intended traffic to be serviced by the service set by configuring the stateful-firewall rules and NAT rules to translate only the traffic from the customer private source address ranges and intended applications. Although this does not prevent unintended traffic from being processed by the services PIC, it prevents the creation of flows, objects, and states

that are not consistent with the expected traffic and are likely to be problematic.

-
- Related Documentation**
- [Configuring Address Pools for Network Address Port Translation \(NAPT\) Overview on page 117](#)

CHAPTER 7

Avoiding IPv4 Exhaustion Using Junos Address Aware Network Addressing and Stateful NAT64

- [Sample IPv6 Transition Scenarios on page 87](#)
- [Configuring Stateful NAT64 on page 89](#)

Sample IPv6 Transition Scenarios

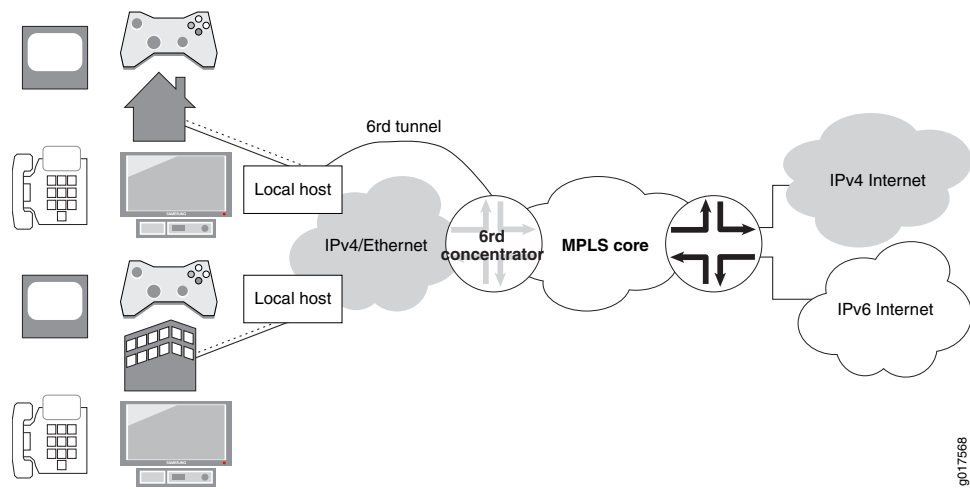
The Junos OS supports many IPv6 transition scenarios required by Junos OS customers. The following are selected examples:

- [Example 1: IPv4 Depletion with a Non-IPv6 Access Network on page 87](#)
- [Example 2: IPv4 Depletion with an IPv6 Access Network on page 88](#)
- [Example 3: IPv4 Depletion for Mobile Networks on page 89](#)

Example 1: IPv4 Depletion with a Non-IPv6 Access Network

[Figure 5 on page 88](#) depicts a scenario in which the Internet service provider (ISP) has not significantly changed its IPv4 network. This approach enables IPv4 hosts to access the IPv4 Internet and IPv6 hosts to access the IPv6 Internet. A dual-stack host can be treated as an IPv4 host when it uses the IPv4 access service, and as an IPv6 host when it uses the IPv6 access service.

Figure 5: IPv4 Depletion Solution - IPv4 Access Network




Two new types of devices must be deployed in this approach: a dual-stack home gateway and a dual-stack carrier-grade Network Address Translation (NAT). The dual-stack home gateway integrates IPv4 forwarding and v6-over-v4 tunneling functions. It can also integrate a v4-v4 NAT function. The dual-stack carrier-grade NAT (CGN) integrates v6-over-v4 tunneling and carrier-grade v4-v4 NAT functions.

Example 2: IPv4 Depletion with an IPv6 Access Network

In the scenario shown in [Figure 6 on page 88](#), the ISP network is IPv6-only.

Figure 6: IPv4 Depletion Solution - IPv6 Access Network



The dual-stack lite (DS-Lite) solution accommodates IPv6-only ISPs. The best business model for this approach is that the customer premises equipment (CPE) has integrated the functions for tunneling IPv6 to an IPv4 backbone, tunneling IPv4 to an IPv6 backbone, and can automatically detect which solution is required.

Not all customers of a given ISP must switch from IPv4 access to IPv6 access simultaneously; in fact, transition can be managed better by switching groups of

customers (for example, all those connected to a single point of presence) on an incremental basis. Such an incremental approach should prove easier to plan, schedule, and execute than an across-the-board conversion.

Example 3: IPv4 Depletion for Mobile Networks

The complexity of mobile networks necessitates a flexible migration approach to ensure minimal disruption and maximum backward compatibility during transition. NAT64 can be used to enable IPv6 devices to communicate to IPv4 hosts without modifying the clients.

Configuring Stateful NAT64

Stateful NAT64 is a mechanism used to move to an IPv6 network and at the same time deal with IPv4 address depletion. By allowing IPv6-only clients to contact IPv4 servers using unicast UDP, TCP, or ICMP, several IPv6-only clients can share the same public IPv4 server address. To allow sharing of the IPv4 server address, stateful NAT64 translates incoming IPv6 packets into IPv4, and vice versa.

To configure stateful NAT64, you must configure a rule at the **[edit services nat]** hierarchy level for translating the source address dynamically and the destination address statically.



BEST PRACTICE: When you configure the service set that includes your NAT rule, include the set `stateful-nat64 clear-dont-fragment-bit` at the **[edit services service-set service-set-name]** hierarchy level. This clears the DF (don't fragment) bit in order to prevent unnecessary creation of an IPv6 fragmentation header when translating IPv4 packets that are less than 1280 bytes. RFC 6145, *IP/ICMP Translation Algorithm*, provides a full discussion of the use of the DF flag to control generation of fragmentation headers. For more information on service sets for NAT, see [“Configuring Service Sets for Network Address Translation” on page 75](#).

To configure stateful NAT64:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

```
[edit]
user@host# edit services nat
```

2. Define the pool of source addresses to be used for dynamic translation.

```
[edit services nat]
user@host# set pool pool name address source addresses
user@host# set pool pool name port source ports
```

For example:

```
[edit services nat]
user@host# set pool src-pool-nat64 address 203.0.113.0/24
user@host# set pool src-pool-nat64 port automatic
```



NOTE: Starting with Junos OS release 14.2, the `sequential` option is introduced to enable you to configure sequential allocation of ports. The `sequential` and `random-allocation` options available with the `port automatic` statement at the `[edit services nat pool nat-pool-name]` hierarchy level are mutually exclusive. You can include the `sequential` option for sequential allocation and the `random-allocation` option for random delegation of ports. By default, sequential allocation of ports takes place if you include only the `port automatic` statement at the `[edit services nat pool nat-pool-name]` hierarchy level. The `auto` option is hidden and is deprecated in Junos OS Release 14.2 and later, and is only maintained for backward compatibility. It might be removed completely in a future software release.

3. Define a NAT rule for translating the source addresses. Set the `match-direction` statement of the rule as `input`. Then define a term that uses `stateful-nat64` as the translation type for translating the addresses of the pool defined in the previous step.

```
[edit services nat]
user@host# set rule rule name match-direction input
user@host# set rule rule name term term name from source-address source address
user@host# set rule rule name term term name from destination-address destination address
user@host# set rule rule name term term name then translated source-pool pool name
user@host# set rule rule name term term name then translated destination-prefix destination prefix
user@host# set rule rule name term term name then translated translation-type stateful-nat64
```

For example:

```
[edit services nat]
user@host# set rule stateful-nat64 match-direction input
user@host# set rule stateful-nat64 term t1 from source-address 2001:DB8::0/96
user@host# set rule stateful-nat64 term t1 from destination-address 64:FF9B::/96
user@host# set rule stateful-nat64 term t1 then translated source-pool src-pool-nat64
user@host# set rule stateful-nat64 term t1 then translated destination-prefix 64:FF9B::/96
user@host# set rule stateful-nat64 term t1 then translated translation-type stateful-nat64
```

The following example configures dynamic source address (IPv6-to-IPv4) and static destination address (IPv6-to-IPv4) translation.

```
[edit services]
user@host# show
nat {
    pool src-pool-nat64 {
        address 203.0.113.0/24;
        port {
            automatic;
        }
    }
    rule stateful-nat64 {
        match-direction input;
        term t1 {
```

```
        from {
            source-address {
                2001:db8::0/96;
            }
            destination-address {
                64:ff9b::/96;
            }
        }
        then {
            translated {
                source-pool src-pool-nat64;
                destination-prefix 64:ff9b::/96;
                translation-type {
                    stateful-nat64;
                }
            }
        }
    }
}
service-set sset-nat64 {
    nat-options {
        stateful-nat64 {
            clear-dont-fragment-bit;
        }
    }
    service-set-options;
    nat-rules stateful-nat64;
    interface-service {
        service-interface ms-0/1/0;
    }
}
```



.....

NOTE: If you configure two NAT64 rules and associate them with the same service set, along with stateful firewall rules, and apply the service set on two VLAN-tagged interfaces, for traffic that is transmitted matching both the NAT rules, the traffic that is destined to the second NAT rule is dropped. In such a scenario, traffic flows are not dropped on the Routing Engine. This behavior of traffic drop by the second NAT rule is expected. With Junos OS Extension-Provider packages installed on a device, because endpoint-independent mapping (EIM) is not supported, EIM per VLAN or per NAT rule term. The second session, which is dropped by the second NAT rule in the configuration scenario described here, is not created owing to the following sequence of events:

1. The first packet matching either rule creates an EIM and a session.
2. The second packet matches the EIM entry because the second packet is sent with the same source IP address and port as the first packet (but with a different destination address).

This condition causes allocation (reuse) of the same public IP address and port to the second packet as the first packet. The reverse flow for this session has the same 5-tuple data as the reverse flow of the first session. This behavior causes flow addition failure because a duplicate flow in the same service set is not permitted.

To work around this problem, disable EIM in both the NAT rules, which causes both the sessions to be established and processed correctly. Alternatively, to avoid this problem, specify the NAT rules on different service-sets configured on different units of the media interface with EIM enabled to successfully establish both the sessions.

.....

CHAPTER 8

Hiding Private Networks Using Static Source NAT

- [Configuring Static Source Translation in IPv4 Networks on page 93](#)
- [Configuring Static Source Translation in IPv6 Networks on page 99](#)
- [Example: Configuring Basic NAT44 on page 103](#)
- [Example: Configuring NAT for Multicast Traffic on page 105](#)

Configuring Static Source Translation in IPv4 Networks

To configure the translation type as **basic-nat44**, you must configure the NAT pool and rule, service set with service interface, and trace options. This topic includes the following tasks:

- [Configuring the NAT Pool and Rule on page 93](#)
- [Configuring the Service Set for NAT on page 95](#)
- [Configuring Trace Options on page 97](#)
- [Sample Configuration - Static Source NAT Using a Static Pool With An Address Prefix And An Address Range on page 98](#)
- [Sample Configuration - Static Source Nat for One-To-One Mapping Between a Private Subnet and a Public Subnet on page 98](#)

Configuring the NAT Pool and Rule

To configure the NAT pool, rule, and term:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

```
[edit]
user@host# edit services nat
```

2. Configure the NAT pool with an address.

```
[edit services nat]
user@host# set pool pool name address address
```

In the following example, the pool name is **src_pool** and the address is **10.10.10.2/32**.

```
[edit services nat]
user@host# set pool src_pool address 10.10.10.2/32
```

3. Configure the NAT rule and the match direction.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction
```

In the following example, the NAT rule name is **rule-basic-nat44** and the match direction is **input**.

```
[edit services nat]
user@host# set rule rule-basic-nat44 match-direction input
```

4. Configure the source address in the **from** statement.

```
[edit services nat]
user@host# set rule rule-basic-nat44 term term-name from from
```

In the following example, the term name is **t1** and the input condition is **source-address 3.1.1.2/32**.

```
[edit services nat]
user@host# set rule rule-basic-nat44 term t1 from source-address 3.1.1.2/32
```

5. Configure the NAT term action and properties of the translated traffic.

```
[edit services nat]
user@host# set rule rule-basic-nat44 term t1 then term-action translated-property
```

In the following example, the term action is **translated** and the property of the translated traffic is **source-pool src_pool**.

```
[edit services nat]
user@host# set rule rule-basic-nat44 term t1 then translated source-pool src_pool
```

6. Configure the translation type.

```
[edit services nat]
user@host# set rule rule-basic-nat44 term t1 then translated translation-type
translation-type
```

In the following example, the translation type is **basic-nat44**.

```
[edit services nat]
user@host# set rule rule-basic-nat44 term t1 then translated translation-type
basic-nat44
```

7. Verify the configuration by using the **show** command at the **[edit services nat]** hierarchy level.

```
[edit services]
user@host# show
nat {
  pool src_pool {
    address 10.10.10.2/32;
  }
  rule rule-basic-nat44 {
    match-direction input;
    term t1 {
      from {
        source-address {
          3.1.1.2/32;
        }
      }
    }
    then {

```



```

        translated {
            source-pool src_pool;
            translation-type {
                basic-nat44;
            }
        }
    }
}

```



NOTE: If you don't configure a stateful firewall (SFW) rule for your traffic, then each packet is subjected to the following default stateful firewall rule:

- Allow any valid packets from inside to outside.
- Create forward and return flow based on packets 5-tuple.
- Allow only valid packets matching return flows from outside to inside.

The stateful firewall's packet validity checks are described in the *Stateful Firewall Anomaly Checking* in “[Junos Network Secure Overview](#)” on page 327. When a packets pass stateful firewall validity checking but are not matched by a NAT rule, they are not translated and may be forwarded if the NAT node has a valid route to the packets' destination IP addresses.



NOTE: When you add or delete a parameter in the from statement (NAT rule term match condition) at the [edit services service-set service-set-name nat-rules rule-name term term-name] hierarchy level, this configuration change triggers a deletion and addition of the NAT policy (which is equivalent to the deactivation and activation of a service set) that causes all existing NAT mappings to be deleted. Because the sessions are not closed owing to the change in the NAT policy, this behavior causes the mappings to timeout immediately after the sessions are closed. This behavior is expected and is applicable only with Junos OS Extension-Provider packages installed on a device. When a NAT policy is deleted and readdded, only EIM mappings are deleted. This NAT policy change does not deactivate and activate the service set. We recommend that you deactivate and reactivate the service set in such scenarios in Junos OS Release 14.2 and earlier.

Configuring the Service Set for NAT

To configure the service set for NAT:

1. In configuration mode, go to the [edit services] hierarchy level.

```

[edit]
user@host# edit services

```

2. Configure the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

In the following example, the service set name is **s1**.

```
[edit services]
user@host# edit service-set s1
```

3. For the **s1** service set, set the reference to the NAT rules configured at the **[edit services nat]** hierarchy level.

```
[edit services service-set s1]
user@host# set nat-rules rule-name
```

In the following example, the rule name is **rule-basic-nat44**.

```
[edit services service-set s1]
user@host# set nat-rules rule-basic-nat44
```

4. Configure the service interface.

```
[edit services service-set s1]
user@host# set interface-service service-interface service-interface-name
```

In the following example, the service interface name is **ms-1/2/0**.

```
[edit services service-set s1]
user@host# set interface-service service-interface ms-1/2/0
```



NOTE: If you have a Trio-based line card, you can configure an inline-services interface on that card:

```
[edit]
user@host# set interfaces si-0/0/0
[edit services service-set s1]
user@host# set interface-service service-interface si-0/0/0
```

5. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
service-set s1 {
    nat-rules rule-basic-nat44;
    interface-service {
        service-interface ms-1/2/0;
    }
}
```

6. Associate the NAT service set with an **xe-** interface:

```
user@host# set interfaces xe-1/1/0 unit 0 family inet address 10.255.247.2/24
user@host# set interfaces xe-1/1/0 unit 0 family inet service input service-set s1
user@host# set interfaces xe-1/1/0 unit 0 family inet service output service-set s1
```

7. Verify the configuration by using the **show** command at the **[edit interfaces]** hierarchy level.

```
[edit interfaces]
user@host# show
```

```

xe-1/1/0 {
  unit 0 {
    family inet {
      service {
        input {
          service-set s1;
        }
        output {
          service-set s1;
        }
      }
      address 10.255.247.2/24;
    }
  }
}

```

Configuring Trace Options

To configure the trace options:

1. In configuration mode, go to the **[edit services adaptive-services-pics]** hierarchy level.

```

[edit]
user@host# edit services adaptive-services-pics

```

2. Configure the trace options.

```

[edit services adaptive-services-pics]
user@host# set traceoptions flag tracing parameter

```

In the following example, the tracing parameter is **all**.

```

[edit services adaptive-services-pics]
user@host# set traceoptions flag all

```

3. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```

[edit services]
user@host# show
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}

```

```

[edit]
user@host# show services
service-set s1 {
  nat-rules rule-basic-nat44;
  interface-service {
    service-interface ms-1/2/0;
  }
}
nat {
  pool src_pool {
    address 10.10.10.2/32;
  }
  rule rule-basic-nat44 {
    match-direction input;
    term t1 {

```

```

        from {
            source-address {
                3.1.1.2/32;
            }
        }
        then {
            translated {
                source-pool src_pool;
                translation-type {
                    basic-nat44;
                }
            }
        }
    }
}

adaptive-services-pics {
    traceoptions {
        flag all;
    }
}

```

Sample Configuration - Static Source NAT Using a Static Pool With An Address Prefix And An Address Range

```
[edit services nat]
pool p1 {
  address 30.30.30.252/30;
  address-range low 20.20.20.1 high 20.20.20.2;
}
rule r1 {
  match-direction input;
  term {
    from {
      source-address {
        10.10.10.252/30;
      }
    }
  }
  then {
    translated {
      source-pool p1;
      translation-type basic-nat44;
    }
  }
}
```

Sample Configuration - Static Source Nat for One-To-One Mapping Between a Private Subnet and a Public Subnet

```
[edit]
user@host# show services
service-set s1 {
    nat-rules rule-basic-nat44;
    interface-service {
        service-interface ms-1/2/0;
    }
}
```

```

}
nat {
  pool src_pool {
    address 10.10.10.2/32;
  }
  rule rule-basic-nat44 {
    match-direction input;
    term t1 {
      from {
        source-address {
          3.1.1.2/32;
        }
      }
      then {
        translated {
          source-pool src_pool;
          translation-type {
            basic-nat44;
          }
        }
      }
    }
  }
}
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}

[edit interfaces]
user@host# show
xe-1/1/0 {
  unit 0 {
    family inet {
      service {
        input {
          service-set s1;
        }
        output {
          service-set s1;
        }
      }
      address 10.255.247.2/24;
    }
  }
}

```

Configuring Static Source Translation in IPv6 Networks

To configure the translation type as **basic-nat66**, you must configure the NAT pool and rule, service set with service interface, and trace options. This topic includes the following tasks:

- [Configuring the NAT Pool and Rule on page 100](#)
- [Configuring the Service Set for NAT on page 101](#)
- [Configuring Trace Options on page 102](#)

Configuring the NAT Pool and Rule

To configure the NAT pool, rule, and term:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

```
[edit]
user@host# edit services nat
```

2. Configure the NAT pool with an address.

```
[edit services nat]
user@host# set pool pool name address address
```

In the following example, the pool name is **src_pool** and the address is **10.10.10.2/32**.

```
[edit services nat]
user@host# set pool src_pool address 10.10.10.2/32
```

3. Configure the NAT rule and the match direction.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction
```

In the following example, the rule name is **rule-basic-nat66** and the match direction is **input**.

```
[edit services nat]
user@host# set rule rule-basic-nat66 match-direction input
```

4. Configure the source address in the **from** statement.

```
[edit services nat]
user@host# set rule rule-basic-nat66 term term-name from from
```

In the following, the term name is **t1** and the input condition is **source-address 10:10:10::0/96**.

```
[edit services nat]
user@host# set rule rule-basic-nat66 term t1 from source-address 10:10:10::0/96
```

5. Configure the NAT term action and properties of the translated traffic.

```
[edit services nat]
user@host# set rule rule-basic-nat66 term t1 then term-action translated-property
```

In the following example, the term action is **translated** and the property of the translated traffic is **source-pool src_pool**.

```
[edit services nat]
user@host# set rule rule-basic-nat66 term t1 then translated source-pool src_pool
```

6. Configure the translation type.

```
[edit services nat]
user@host# set rule rule-basic-nat66 term t1 then translated translation-type
translation-type
```

In the following example, the translation type is **basic-nat66**.

```
[edit services nat]
```

```
user@host# set rule rule-basic-nat66 term t1 then translated translation-type
basic-nat66
```

7. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
nat {
  pool src_pool {
    address 10.10.10.2/32;
  }
  rule rule-basic-nat66 {
    match-direction input;
    term t1 {
      from {
        source-address {
          10:10:10::0/96;
        }
      }
      then {
        translated {
          source-pool src_pool;
          translation-type {
            basic-nat66;
          }
        }
      }
    }
  }
}
```

Configuring the Service Set for NAT

To configure the service set for NAT:

1. In configuration mode, go to the **[edit services]** hierarchy level.

```
[edit]
user@host# edit services
```

2. Configure the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

In the following example, the service set name is **s1**.

```
[edit services]
user@host# edit service-set s1
```

3. For the **s1** service set, set the reference to the NAT rules configured at the **[edit services nat]** hierarchy level.

```
[edit services service-set s1]
user@host# set nat-rules rule-name
```

In the following example, the rule name is **rule-basic-nat66**.

```
[edit services service-set s1]
user@host# set nat-rules rule-basic-nat66
```

4. Configure the service interface.

```
[edit services service-set s1]
user@host# set interface-service service-interface service-interface-name
```

In the following example, the service interface name is **sp-1/2/0**.

```
[edit services service-set s1]
user@host# set interface-service service-interface sp-1/2/0
```

5. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
service-set s1 {
  nat-rules rule-basic-nat66;
  interface-service {
    service-interface sp-1/2/0;
  }
}
```

Configuring Trace Options

To configure the trace options at the **[edit services adaptive-services-pics]** hierarchy level:

1. In configuration mode, go to the **[edit services adaptive-services-pics]** hierarchy level.

```
[edit]
user@host# edit services adaptive-services-pics
```

2. Configure the trace options.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag tracing parameter
```

In the following example, the tracing parameter is **all**.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag all
```

3. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```

The following example configures the translation type as **basic-nat66**.

```
[edit]
user@host# show services
service-set s1 {
  nat-rules rule-basic-nat66;
  interface-service {
    service-interface sp-1/2/0;
  }
}
```



```

}
nat {
  pool src_pool {
    address 10.10.10.2/32;
  }
  rule rule-basic-nat66 {
    match-direction input;
    term t1 {
      from {
        source-address {
          10:10:10::0/96;
        }
      }
      then {
        translated {
          source-pool src_pool;
          translation-type {
            basic-nat66;
          }
        }
      }
    }
  }
}
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}

```

Example: Configuring Basic NAT44

This example describes how to implement a basic NAT44 configuration.

- [Requirements on page 103](#)
- [Overview on page 103](#)
- [Configuring Basic NAT44 on page 104](#)

Requirements

This example uses the following hardware and software components:

- An MX Series 3D Universal Edge router with a Services DPC or an M Series Multiservice Edge router with a services PIC
- A domain name server (DNS)
- Junos OS Release 11.4 or higher

Overview

This example shows a complete CGN NAT44 configuration and advanced options.

Configuring Basic NAT44

Chassis Configuration

Step-by-Step Procedure

To configure the service PIC (FPC 5 Slot 0) with the Layer 3 service package:

1. Go to the **[edit chassis]** hierarchy level.

```
user@host# edit chassis
```
2. Configure the Layer 3 service package.

```
[edit chassis]  
user@host# set fpc 5 pic 0 adaptive-services service-package layer-3
```

Interfaces Configuration

Step-by-Step Procedure

To configure interfaces to the private network and the public Internet:

1. Define the interface to the private network.

```
user@host# edit interfaces ge-1/3/5  
[edit interfaces ge-1/3/5]  
user@host# set description "Private"  
user@host# edit unit 0 family inet  
[edit interfaces ge-1/3/5 unit 0 family inet]  
user@host# set service input service-set ss2  
user@host# set service output service-set ss2  
user@host# set address 9.0.0.1/24
```
2. Define the interface to the public Internet.

```
user@host# edit interfaces ge-1/3/6  
[edit interfaces ge-1/3/6]  
user@host# set description "Public"  
user@host# set unit 0 family inet address 128.0.0.1/24
```
3. Define the service interface for NAT processing.

```
user@host# edit interfaces sp-5/0/0  
[edit interfaces sp-5/0/0]  
user@host# set unit 0 family inet
```

```

Results user@host# show interfaces ge-1/3/5
description Private;
unit 0 {
  family inet {
    service {
      input {
        service-set sset2;
      }
      output {
        service-set sset2;
      }
    }
    address 9.0.0.1/24;
  }
}

user@host# show interfaces ge-1/3/6
description Public;;
unit 0 {
  family inet {
    address 128.0.0.1/24;
  }
}


user@host# show interfaces sp-5/0/0
unit 0 {
  family inet;
}

```

Example: Configuring NAT for Multicast Traffic

Figure 7 on page 105 illustrates the network setup for the following configuration, which allows IP multicast traffic to be sent to the Multiservices PIC.

Figure 7: Configuring NAT for Multicast Traffic



- [Rendezvous Point Configuration on page 105](#)
- [Router 1 Configuration on page 108](#)

Rendezvous Point Configuration

On the rendezvous point (RP), all incoming traffic from the multicast source at **192.168.254.0/27** is sent to the static NAT pool **mcast_pool**, where its source is translated to **20.20.20.0/27**. The service set **nat_ss** is a next-hop service set that allows IP multicast traffic to be sent to the Multiservices DPC or Multiservices PIC. The inside interface on the PIC is **ms-1/1/0.1** and the outside interface is **ms-1/1/0.2**.

[edit services]

```
nat {
  pool mcast_pool {
    address 20.20.20.0/27;
  }
  rule nat_rule_1 {
    match-direction input;
    term 1 {
      from {
        source-address 192.168.254.0/27;
      }
    }
    then {
      translated {
        source-pool mcast_pool;
        translation-type basic-nat44;
      }
      syslog;
    }
  }
}
service-set nat_ss {
  allow-multicast;
  nat-rules nat_rule_1;
  next-hop-service {
    inside-service-interface ms-1/1/0.1;
    outside-service-interface ms-1/1/0.2;
  }
}
```

The Gigabit Ethernet interface **ge-0/3/0** carries traffic out of the RP to Router 1. The multiservices interface **ms-1/1/0** has two logical interfaces: **unit 1** is the inside interface for next-hop services and **unit 2** is the outside interface for next-hop services. Multicast source traffic comes in on the Fast Ethernet interface **fe-1/2/1**, which has the firewall filter **fbf** applied to incoming traffic.

```
[edit interfaces]
ge-0/3/0 {
  unit 0 {
    family inet {
      address 10.10.1.1/30;
    }
  }
}
ms-1/1/0 {
  unit 0 {
    family inet;
  }
  unit 1 {
    family inet;
    service-domain inside;
  }
  unit 2 {
    family inet;
    service-domain outside;
  }
}
```

```

}
fe-1/2/1 {
  unit 0 {
    family inet {
      filter {
        input fbf;
      }
      address 192.168.254.27/27;
    }
  }
}

```

Multicast packets can only be directed to the Multiservices DPC or the Multiservices PIC using a next-hop service set. In the case of NAT, you must also configure a VPN routing and forwarding instance (VRF). Therefore, the routing instance **stage** is created as a “dummy” forwarding instance. To direct incoming packets to **stage**, you configure filter-based forwarding through a firewall filter called **fbf**, which is applied to the incoming interface **fe-1/2/1**. A lookup is performed in **stage.inet.0**, which has a multicast static route that is installed with the next hop pointing to the PIC's inside interface. All multicast traffic matching this route is sent to the PIC.

```

[edit firewall]
filter fbf {
  term 1 {
    then {
      routing-instance stage;
    }
  }
}

```

The routing instance **stage** forwards IP multicast traffic to the inside interface **ms-1/1/0.1** on the Multiservices DPC or Multiservices PIC:

```

[edit]
routing-instances stage {
  instance-type forwarding;
  routing-options {
    static {
      route 224.0.0.0/4 next-hop ms-1/1/0.1;
    }
  }
}

```

You enable OSPF and Protocol Independent Multicast (PIM) on the Fast Ethernet and Gigabit Ethernet logical interfaces over which IP multicast traffic enters and leaves the RP. You also enable PIM on the outside interface (**ms-1/1/0.2**) of the next-hop service set.

```

[edit protocols]
ospf {
  area 0.0.0.0 {
    interface fe-1/2/1.0 {
      passive;
    }
    interface lo0.0;
    interface ge-0/3/0.0;
  }
}

```

```
    }  
  }  
  pim {  
    rp {  
      local {  
        address 10.255.14.160;  
      }  
    }  
    interface fe-1/2/1.0;  
    interface lo0.0;  
    interface ge-0/3/0.0;  
    interface ms-1/1/0.2;  
  }
```

As with any filter-based forwarding configuration, in order for the static route in the forwarding instance **stage** to have a reachable next hop, you must configure routing table groups so that all interface routes are copied from **inet.0** to the routing table in the forwarding instance. You configure routing tables **inet.0** and **stage.inet.0** as members of **fbf_rib_group**, so that all interface routes are imported into both tables.

```
[edit routing-options]  
interface-routes {  
  rib-group inet fbf_rib_group;  
}  
rib-groups fbf_rib_group {  
  import-rib [ inet.0 stage.inet.0 ];  
}  
multicast {  
  rpf-check-policy no_rpf;  
}
```

Reverse path forwarding (RPF) checking must be disabled for the multicast group on which source NAT is applied. You can disable RPF checking for specific multicast groups by configuring a policy similar to the one in the example that follows. In this case, the **no_rpf** policy disables RPF check for multicast groups belonging to **224.0.0.0/4**.

```
[edit policy-options]  
policy-statement no_rpf {  
  term 1 {  
    from {  
      route-filter 224.0.0.0/4 orlonger;  
    }  
    then reject;  
  }  
}
```

Router 1 Configuration

The Internet Group Management Protocol (IGMP), OSPF, and PIM configuration on Router 1 is as follows. Because of IGMP static group configuration, traffic is forwarded out **fe-3/0/0.0** to the multicast receiver without receiving membership reports from host members.

```
[edit protocols]  
igmp {
```

```
interface fe-3/0/0.0 {  
}  
}  
ospf {  
  area 0.0.0.0 {  
    interface fe-3/0/0.0 {  
      passive;  
    }  
    interface lo0.0;  
    interface ge-7/2/0.0;  
  }  
  pim {  
    rp {  
      static {  
        address 10.255.14.160;  
      }  
    }  
    interface fe-3/0/0.0;  
    interface lo0.0;  
    interface ge-7/2/0.0;  
  }  
}
```

The routing option creates a static route to the NAT pool, **mcast_pool**, on the RP.

```
[edit routing-options]  
static {  
  route 20.20.20.0/27 next-hop 10.10.1.1;  
}
```


CHAPTER 9

Making Private Servers Available Using Static Destination NAT

- [Configuring Static Destination Address Translation in IPv4 Networks on page 111](#)

Configuring Static Destination Address Translation in IPv4 Networks

In IPv4 networks, destination address translation is a mechanism used to implement address translation for destination traffic without port mapping. To use destination address translation, the size of the pool address space must be greater than or equal to the destination address space. You must specify a name for the **destination-pool** statement, which can contain multiple addresses, ranges, or prefixes, as long as the number of NAT addresses in the pool is larger than the number of destination addresses in the **from** statement.

To configure destination address translation in IPv4 networks:

1. In configuration mode, go to the **[edit services]** hierarchy level.

```
[edit]
user@host# edit services
```

2. Configure the service set and the NAT rule.

```
[edit services]
user@host# set service-set service-set-name nat-rules rule-name
```

In the following example, the name of the service set is **s1** and the name of the NAT rule is **rule-dnat44**.

```
[edit services]
user@host# set service-set s1 nat-rules rule-dnat44
```

3. Go to the **[interface-service]** hierarchy level of the service set.

```
[edit services]
user@host# edit service-set s1 interface-service
```

4. Configure the service interface.

```
[edit services service-set s1 interface-service]
user@host# set service-interface service-interface-name
```

In the following example, the name of the service interface is **ms-0/1/0**.



NOTE: If the service interface is not present in the router, or the specified interface is not functional, the following command can result in an error.

```
[edit services service-set s1 interface-service]
user@host# set service-interface ms-0/1/0
```

5. Go to the **[edit services nat]** hierarchy level. Issue the following command from the top of the services hierarchy, or use the **top** keyword.

```
[edit services service-set s1]
user@host# top edit services nat
```

6. Configure the NAT pool with an address.

```
[edit services nat]
user@host# set pool pool-name address address
```

In the following example, **dest-pool** is used as the pool name and **4.1.1.2** as the address.

```
user@host# set pool dest-pool address 4.1.1.2
```

7. Configure the rule, match direction, term, and destination address.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction term term-name from
destination-address address
```

In the following example, the name of the rule is **rule-dnat44**, the match direction is **input**, the name of the term is **t1**, and the address is **20.20.20.20**.

```
[edit services nat]
user@host# set rule rule-dnat44 match-direction input term t1 from
destination-address 20.20.20.20
```

8. Go to the **[edit services nat rule rule-dnat44 term t1]** hierarchy level.

```
[edit services nat]
user@host# edit rule rule-dnat44 term t1
```

9. Configure the destination pool and the translation type.

```
[edit services nat rule rule-dnat44 term t1]
user@host# set then translated destination-pool dest-pool-name translation-type
translation-type
```

In the following example, the destination pool name is **dest-pool**, and the translation type is **dnat-44**.

```
[edit services nat rule rule-dnat44 term t1]
user@host# set then translated destination-pool dest-pool translation-type dnat-44
```

10. Go to the **[edit services adaptive-services-pics]** hierarchy level. In the following command, the **top** keyword ensures that the command is run from the top of the hierarchy.

```
[edit services nat rule rule-dnat44 term t1]
user@host# top edit services adaptive-services-pics
```

11. Configure the trace options.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag tracing parameter
```

In the following example, the tracing parameter is configured as **all**.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag all
```

12. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
service-set s1 {
  nat-rules rule-dnat44;
  interface-service {
    service-interface ms-0/1/0;
  }
}
nat {
  pool dest-pool {
    address 4.1.1.2/32;
  }
  rule rule-dnat44 {
    match-direction input;
    term t1 {
      from {
        destination-address {
          20.20.20.20/32;
        }
      }
      then {
        translated {
          destination-pool dest-pool;
          translation-type {
            dnat-44;
          }
        }
      }
    }
  }
}
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```

The following example configures the translation type as **dnat-44**.

```
[edit services]
user@host# show
service-set s1 {
  nat-rules rule-dnat44;
  interface-service {
    service-interface ms-0/1/0;
  }
}
nat {
  pool dest-pool {
    address 4.1.1.2/32;
  }
}
```

```

rule rule-dnat44 {
    match-direction input;
    term t1 {
        from {
            destination-address {
                20.20.20.20/32;
            }
        }
        then {
            translated {
                destination-pool dest-pool;
                translation-type {
                    dnat-44;
                }
            }
        }
    }
}

adaptive-services-pics {
    traceoptions {
        flag all;
    }
}

```

In the following configuration, **term1** configures source address translation for traffic from any private address to any public address. The translation is applied for all services. **term2** performs destination address translation for Hypertext Transfer Protocol (HTTP) traffic from any public address to the server's virtual IP address. The virtual server IP address is translated to an internal IP address.

```
[edit services nat]
rule my-nat-rule {
  match-direction input;
  term my-term1 {
    from {
      source-address private;
      destination-address public;
    }
    then {
      translated {
        source-pool my-pool; # pick address from a pool
        translation-type napt-44; # dynamic NAT with port translation
      }
    }
  }
}

rule my-nat-rule2 {
  match-direction input;
  term my-term2 {
    from {
      destination-address 192.168.137.3; # my server's virtual address
      application http;
    }
    then {
      translated {
        destination-pool nat-pool-name;
        translation-type dnat-44; # static destination NAT
      }
    }
  }
}
```

```
    }  
  }  
}  
}
```

The following configuration performs NAT using the destination prefix **20.20.10.0/32** without defining a pool.

```
[edit services nat]  
rule src-nat {  
  match-direction input;  
  term t1 {  
    from {  
      destination-address 10.10.10.10/32;  
      then {  
        translation-type dnat44;  
        destination-prefix 20.20.10.0/24;  
      }  
    }  
  }  
}
```


CHAPTER 10

Allowing Components of a Private Network to Share a Single Address Using NAT

- [Configuring Address Pools for Network Address Port Translation \(NAPT\) Overview on page 117](#)
- [Configuring Dynamic Source Address and Port Translation in IPv4 Networks on page 127](#)
- [Configuring Dynamic Source Address and Port Translation for IPv6 Networks on page 131](#)
- [Example: Configuring NAT with Port Translation on page 133](#)
- [Example: NAPT Configuration for the MS-MPC on page 134](#)
- [Example: Dynamic Source NAT as a Next-Hop Service on page 138](#)

Configuring Address Pools for Network Address Port Translation (NAPT) Overview

With Network Address Port Translation (NAPT), you can configure up to 32 address ranges with up to 65,536 addresses each.

The **port** statement specifies port assignment for the translated addresses. To configure automatic assignment of ports, include the **port automatic** statement at the **[edit services nat pool nat-pool-name]** hierarchy level. By default, sequential allocation of ports occurs. You can include the **sequential** option with the **port automatic** statement at the **[edit services nat pool nat-pool-name]** hierarchy level, starting with Junos OS Release 14.2 for sequenced allocation of ports from the specified range. To configure a specific range of port numbers, include the **port range low minimum-value high maximum-value** statement at the **[edit services nat pool nat-pool-name]** hierarchy level.



NOTE: When 99% of the total available ports in pool for napt-44 , no new flows are allowed on that NAT pool.

The **auto** option is hidden and is deprecated in Junos OS Release 14.2 and later, and is only maintained for backward compatibility. It might be removed completely in a future software release.

The Junos OS provides several alternatives for allocating ports:

- [Round-Robin Allocation for NAPT on page 118](#)
- [Sequential Allocation for NAPT on page 118](#)
- [Preserve Parity and Preserve Range for NAPT on page 119](#)
- [Address Pooling and Endpoint Independent Mapping for NAPT on page 119](#)
- [Port Block Allocation for NAPT on page 121](#)
- [Deterministic Port Block Allocation for NAPT on page 122](#)
- [Comparison of NAPT Implementation Methods on page 126](#)

Round-Robin Allocation for NAPT

To configure round-robin allocation for NAT pools, include the **address-allocation round-robin** configuration statement at the **[edit services nat pool *pool-name*]** hierarchy level. When you use round-robin allocation, one port is allocated from each address in a range before repeating the process for each address in the next range. After ports have been allocated for all addresses in the last range, the allocation process wraps around and allocates the next unused port for addresses in the first range.

- The first connection is allocated to the address:port 100.0.0.1:3333.
- The second connection is allocated to the address:port 100.0.0.2:3333.
- The third connection is allocated to the address:port 100.0.0.3:3333.
- The fourth connection is allocated to the address:port 100.0.0.4:3333.
- The fifth connection is allocated to the address:port 100.0.0.5:3333.
- The sixth connection is allocated to the address:port 100.0.0.6:3333.
- The seventh connection is allocated to the address:port 100.0.0.7:3333.
- The eighth connection is allocated to the address:port 100.0.0.8:3333.
- The ninth connection is allocated to the address:port 100.0.0.9:3333.
- The tenth connection is allocated to the address:port 100.0.0.10:3333.
- The eleventh connection is allocated to the address:port 100.0.0.11:3333.
- The twelfth connection is allocated to the address:port 100.0.0.12:3333.
- Wraparound occurs and the thirteenth connection is allocated to the address:port 100.0.0.1:3334.

Sequential Allocation for NAPT

With sequential allocation, the next available address in the NAT pool is selected only when all the ports available from an address are exhausted.



NOTE: This legacy implementation provides backward compatibility and is no longer a recommended approach.

The NAT pool called **napt** in the following configuration example uses the sequential implementation:

```
pool napt {  
  address-range low 100.0.0.1 high 100.0.0.3;  
  address-range low 100.0.0.4 high 100.0.0.6;  
  address-range low 100.0.0.8 high 100.0.0.10;  
  address-range low 100.0.0.12 high 100.0.0.13;  
  port {  
    range low 3333 high 3334;  
  }  
}
```

In this example, the ports are allocated starting from the first address in the first address-range, and allocation continues from this address until all available ports have been used. When all available ports have been used, the next address (in the same address-range or in the following address-range) is allocated and all its ports are selected as needed. In the case of the example **napt** pool, the tuple address, port 100.0.0.4:3333, is allocated only when all ports for all the addresses in the first range have been used.

- The first connection is allocated to the address:port 100.0.0.1:3333.
- The second connection is allocated to the address:port 100.0.0.1:3334.
- The third connection is allocated to the address:port 100.0.0.2:3333.
- The fourth connection is allocated to the address:port 100.0.0.2:3334, and so on.

Preserve Parity and Preserve Range for NAPT

The following options are available for NAPT:

- Preserving parity—Use the **preserve-parity** command to allocate even ports for packets with even source ports and odd ports for packets with odd source ports.
- Preserving range—Use the **preserve-range** command to allocate ports within a range from 0 to 1023, assuming the original packet contains a source port in the reserved range. This applies to control sessions, not data sessions.

Address Pooling and Endpoint Independent Mapping for NAPT

- [Address Pooling on page 119](#)
- [Endpoint Independent Mapping and Endpoint Independent Filtering on page 120](#)

Address Pooling

Address pooling, or address pooling paired (APP) ensures assignment of the same external IP address for all sessions originating from the same internal host. You can use this feature when assigning external IP addresses from a pool. This option does not affect port utilization

Address pooling solves the problems of an application opening multiple connections. For example, when Session Initiation Protocol (SIP) client sends Real-Time Transport Protocol (RTP) and Real-Time Control Protocol (RTCP) packets, the SIP generally server

requires that they come from the same IP address, even if they have been subject to NAT. If RTP and RTCP IP addresses are different, the receiving endpoint might drop packets. Any point-to-point (P2P) protocol that negotiates ports (assuming address stability) benefits from address pooling paired.

The following are use cases for address pooling:

- A site that offers instant messaging services requires that chat and their control sessions come from the same public source address. When the user signs on to chat, a control session authenticates the user. A different session begins when the user starts a chat session. If the chat session originates from a source address that is different from the authentication session, the instant messaging server rejects the chat session, because it originates from an unauthorized address.
- Certain websites such as online banking sites require that all connections from a given host come from the same IP address.



.....

NOTE: Starting with Junos OS Release 14.1, when you deactivate a service-set that contains address pooling paired (APP) for that service-set, messages are displayed on the PIC console and the mappings are cleared for that service-set. These messages are triggered when the deletion of a service-set commences and again generated when the deletion of the service-set is completed. The following sample messages are displayed when deletion starts and ends:

- Nov 15 08:33:13.974 LOG: Critical] SVC-SET ss1 (iid 5) deactivate/delete: NAT Mappings and flows deletion initiated
- Nov 15 08:33:14.674 LOG: Critical] SVC-SET ss1 (iid 5) deactivate/delete: NAT Mappings and flows deletion completed

In a scaled environment that contains a large number of APP in a service set, a heavy volume of messages is generated and this process takes some amount of time. We recommend that you wait until the console messages indicating the completion of deletion of the service set are completed before you reactivate the service-set again.

.....

Endpoint Independent Mapping and Endpoint Independent Filtering

Endpoint independent mapping (EIM) ensures the assignment of the same external address *and* port for all connections from a given host if they use the same internal port. This means if they come from a different source port, you are free to assign a different external address.

EIM and APP differ as follows:

- APP ensures assigning the same external IP address.
- EIM provides a stable external IP address and port (for a period of time) to which external hosts can connect. Endpoint independent filtering (EIF) controls which external hosts can connect to an internal host.



NOTE: Starting with Junos OS Release 14.1, when you deactivate a service-set that contains endpoint independent mapping (EIM) mapping for that service-set, messages are displayed on the PIC console and the mappings are cleared for that service-set. These messages are triggered when the deletion of a service-set commences and again generated when the deletion of the service-set is completed. The following sample messages are displayed when deletion starts and ends:

- Nov 15 08:33:13.974 LOG: Critical] SVC-SET ss1 (iid 5) deactivate/delete: NAT Mappings and flows deletion initiated
- Nov 15 08:33:14.674 LOG: Critical] SVC-SET ss1 (iid 5) deactivate/delete: NAT Mappings and flows deletion completed

In a scaled environment that contains a large number of EIM mappings in a service set, a heavy volume of messages is generated and this process takes some amount of time. We recommend that you wait until the console messages indicating the completion of deletion of the service set are completed before you reactivate the service-set again.

Port Block Allocation for NAPT

Carriers track subscribers using the IP address (RADIUS or DHCP) log. If they use NAPT, an IP address is shared by multiple subscribers, and the carrier must track the IP address and port, which are part of the NAT log. Because ports are used and reused at a very high rate, tracking subscribers using the log becomes difficult due to the large number of messages, which are difficult to archive and correlate. By enabling the allocation of ports in blocks, port block allocation can significantly reduce the number of logs, making it easier to track subscribers.

Port block allocation is supported on MX series routers with MultiServices Dense Port Concentrators (MS-DPCs).

- [Secured Port Block Allocation for NAPT on page 121](#)
- [Interim Logging for Port Block Allocation on page 122](#)

Secured Port Block Allocation for NAPT

Secured port block allocation can be used for translation types **napt-44** and **stateful-nat64**.

When allocating blocks of ports, the most recently allocated block is the current active block. New requests for NAT ports are served from the active block. Ports are allocated randomly from the current active block.

When you configure secured port block allocation, you can specify the following:

- **block-size**
- **max-blocks-per-address**

- **active-block-timeout**

Interim Logging for Port Block Allocation

With port block allocation we generate one syslog log per set of ports allocated for a subscriber. These logs are UDP based and can be lost in the network, particularly for long-running flows. Interim logging triggers re-sending the above logs at a configured interval for active blocks that have traffic on at least one of the ports of the block.

Interim logging is activated by including the **pba-interim-logging-interval** statement under **services-options** for sp- interfaces.

Deterministic Port Block Allocation for NAT

You can configure NAT algorithm-based allocation of blocks of destination ports. By specifying **deterministic-port-block-allocation blocksize blocksize** at the **[edit services nat pool poolname port]** hierarchy level, you ensure that an incoming (source) IP address and port always map to the same destination IP address and port, thus eliminating the need for the address translation logging. You can also specify **include-boundary-addresses** if you want the lowest and highest addresses in the source address range of a NAT rule to be translated when the NAT pool is used. When you use deterministic port block allocation, you must specify **deterministic-nat44** as the **translation-type** in your NAT rule.

For detailed information on how to configure deterministic port block allocation, see [“Configuring Deterministic Port Block Allocation” on page 180](#).

- [Understanding Deterministic Port Block Allocation Algorithms on page 122](#)
- [Deterministic Port Block Allocation Algorithm Usage on page 123](#)
- [Deterministic Port Block Allocation Restrictions on page 125](#)

Understanding Deterministic Port Block Allocation Algorithms

The effectiveness of your implementation of deterministic port block allocation depends on your analysis of your subscriber requirements. The block size you provide indicates how many ports will be made available for each incoming subscriber address in the range the **from** clause specified in the applicable NAT rule. The allocation algorithm computes an offset value to determine the outgoing port. A reverse algorithm is used to derive the originating subscriber address.



NOTE: In order to track subscribers without using logs, an ISP must use a reverse algorithm to derive a subscriber (source) addresses from translated addresses.

Deterministic Port Block Allocation Algorithm Usage

When you have configured deterministic port block allocation, you can use the ***show services nat deterministic-nat internal-host*** and ***show services nat deterministic-nat nat-port-block*** commands to show forward and reverse mapping. However, mappings will change if you reconfigure your deterministic port block allocation block size or the **from** clause for your NAT rule. In order to provide historical information on mappings, we recommend that you write scripts that can show specific mappings for prior configurations.

The following variables are used in forward calculation (private subscriber IP address to public IP address) and reverse calculation (public IP address to private subscriber IP address):

- Pr_Prefix—Any pre-NAT IPv4 subscriber address
- Pr_Port—Any pre-NAT protocol port
- Block_Size—Number of ports configured to be available for each Pr_Prefix
- Base_PR_Prefix—First usable pre-NAT IPv4 subscriber address in a “from” clause match condition
- Base_PU_Prefix—First usable post-NAT IPv4 subscriber address configured in the NAT pool.
- Pu_Port_Range_Start—1024 (ports 0 through 1023 are not used when **port automatic** is configured)
- Pr_Offset—Pr_Prefix – Base_Pr_Prefix
- PR_Port_Offset—Pr_Offset * Block_Size
- Pu_Prefix—Post-NAT address for a given Pr_Prefix
- Pu_Start_Port—Post-NAT start port for a flow from a given Pr_Prefix
- Pu_Actual_Port—Post-NAT port seen on a reverse flow
- Nr_Addr_PR_Prefix — Number of usable pre-NAT IPv4 subscriber addresses in a “from” clause match condition
- Nr_Addr_PU_Prefix — Number of usable post-NAT IPv4 addresses configured in the NAT pool
- Rounded_Port_Range_Per_IP — $\text{ceil}[(\text{Nr_Addr_PR_Prefix}/\text{Nr_Addr_PU_Prefix})] * \text{Block_Size}$
- Pu_Offset—Pu_Prefix – Base_Pu_Prefix
- Pu_Port_Offset—(Pu_Offset * Port_Range_Per_Pu_IP) + (Pu_Actual_Port – Pu_Port_Start_Port)



NOTE: If `block-size` is configured as zero, the method for computing the block size has changed and is computed as follows:

$$\text{block-size} = \text{int}(\text{ceil}[(\text{Nr_Addr_PR_Prefix} / \text{Nr_Addr_PU_Prefix})])$$

where 64512 is the maximum available port range per public IP address.

Algorithm Usage—Assume the following configuration:

```
services {
  nat {
    pool src-pool {
      address-range low 32.32.32.1 high 32.32.32.254;
      port {
        automatic {
          random-allocation;
        }
        deterministic-block-allocation {
          block-size 249;
        }
      }
    }
  }
  rule det-nat {
    match-direction input;
    term t1 {
      from {
        source-address {
          10.1.0.0/16;
        }
      }
      then {
        translated {
          source-pool src-pool;
          translation-type {
            deterministic-napt44;
          }
        }
      }
    }
  }
}
```

Forward Translation

1. $\text{Pr_Offset} = \text{Pr_Prefix} - \text{Base_Pr_Prefix}$
2. $\text{Pr_Port_Offset} = \text{Pr_Offset} * \text{Block_Size}$
3. $\text{Rounded_Port_Range_Per_IP} = \text{ceil}[(\text{Nr_Addr_PR_Prefix} / \text{Nr_Addr_PU_Prefix})] * \text{Block_Size}$
4. $\text{Pu_Prefix} = \text{Base_Public_Prefix} + \text{floor}(\text{Pr_Port_Offset} / \text{Rounded_Port_Range_Per_IP})$
5. $\text{Pu_Start_Port} = \text{Pu_Port_Range_Start} + (\text{Pr_Port_Offset} \% \text{Rounded_Port_Range_Per_IP})$

Using the sample configuration and assuming a subscriber flow sourced from 10.1.1.250:5000:

1. $\text{Pr_Offset} = 10.1.1.250 - 10.1.0.1 = 505$
2. $\text{Pu_Port_Offset} = 505 * 249 = 125,745$
3. $\text{Rounded_Port_Range_Per_IP} = \text{ceil}[(65,533/254)] * 249 = 259 * 249 = 64,491$
4. $\text{Pu_Prefix} = 32.32.32.1 + \text{floor}(125,745 / 64,491) = 32.32.32.1 + 1 = 32.32.32.2$
5. $\text{Pu_Start_Port} = 1,024 + (125,745 \% 64,491) = 62278$
 - 10.1.1.250 is translated to 32.32.32.2.
 - The starting port is 62278. There are 249 ports available to the subscriber based on the configured block size. The available port range spans ports 62278 through 62526 (inclusive).
 - The specific flow 10.1.1.250:5000 randomly assigns any of the ports in its range because random allocation was specified.

Reverse Translation

1. $\text{Pu_Offset} = \text{Pu_Prefix} - \text{Base_Pu_Prefix}$
2. $\text{Pu_Port_Offset} = (\text{Pu_Offset} * \text{Rounded_Port_Range_Per_IP}) + (\text{Pu_Actual_Port} - \text{Pu_Port_Range_Start})$
3. $\text{Subscriber_IP} = \text{Base_Pr_Prefix} + \text{floor}(\text{Pu_Port_Offset} / \text{Block_Size})$

The reverse translation is determined as follows. Assume a flow returning to 32.32.32.2:62278.

1. $\text{Pu_Offset} = 32.32.32.2 - 32.32.32.1 = 1$
2. $\text{Pu_Port_Offset} = (1 * 64,491) + (62,280 - 1024) = 125,747$
3. $\text{Subscriber_IP} = 10.1.0.1 + \text{floor}(125,747 / 249) = 10.1.0.1 + 505 = 10.1.1.250$



NOTE: In reverse translation, only the original private IP address can be derived, and not the original port in use. This is sufficiently granular for law enforcement requirements.

Deterministic Port Block Allocation Restrictions

When you configure deterministic port block allocation, you must be aware of the following restrictions. Violation of any restriction results in a commit error. The restrictions and their error messages are shown in [Table 10 on page 126](#).

Table 10: Deterministic Port Block Allocation Commit Constraints

Restriction	Error Message
The total number of deterministic NAT blocks must be greater than or equal to the 'from' clause addresses configured. This means that the Rounded_Port_Range_Per_IP value must be less than or equal to 64,512.	Number of addresses and port blocks combination in the NAT pool is less than number of addresses in 'from' clause
IPv6 addresses should not be used in deterministic NAT pool/from clause.	Invalid IP address in pool p1 with translation type deterministic-napt44 OR There is already a range configured with v4 address range
The from clause addresses should be same if the same deterministic NAT pool is used across multiple terms/rules. Only one from clause address/range should be specified if the same deterministic NAT pool is used across multiple terms/rules.	With translation-type deterministic-napt44, same 'from' address/range should be configured if pool is shared by multiple rules or terms
There shouldn't be address overlap between except entries in the from clause addresses.	overlapping address, in the 'from' clause between 'except' entries
A deterministic NAT pool cannot be used with other translation-types	Deterministic NAT pool cannot be used with other translation-types
Deterministic NAPT44 must use a source pool with deterministic-port-block-allocation configuration	Deterministic NAPT44 must use a source pool with deterministic-port-block-allocation configuration
If address-allocation round-robin is configured, a commit results in display of a warning indicating that this technique is not needed with translation-type deterministic-napt44 and is ignored.	Address allocation round-robin is not needed with translation-type deterministic-napt44
The total number of IP addresses assigned to a deterministic NAT pool should be less than or equal to 2^{24} (16777216).	Number of addresses in pool with deterministic-napt44 translation are limited to at most 16777216 (2^{24})

Comparison of NAPT Implementation Methods

Table 11 on page 126 provides a feature comparison of available NAPT implementation methods.

Table 11: Comparison of NAPT Implementation Methods

Feature/Function	Dynamic Port Allocation	Secured Port Block Allocation	Deterministic Port Block Allocation
Users per IP	High	Medium	Low
Security Risk	Low	Medium	Medium
Log Utilization	High	Low	None (no logs necessary)
Security Risk Reduction	Random allocation	active-block-timeout feature	n/a

Table 11: Comparison of NAT Implementation Methods (*continued*)

Feature/Function	Dynamic Port Allocation	Secured Port Block Allocation	Deterministic Port Block Allocation
Increasing Users per IP	n/a	Configure multiples of smaller port blocks to maximize users/ public IP	Algorithm-based port allocation

Configuring Dynamic Source Address and Port Translation in IPv4 Networks

Network Address Port Translation (NAPT) is a method by which many network addresses and their TCP/UDP ports are translated into a single network address and its TCP/UDP ports. This translation can be configured in both IPv4 and IPv6 networks. This section describes the steps for configuring NAPT in IPv4 networks.

To configure NAPT, you must configure a rule at the **[edit services nat]** hierarchy level for dynamically translating the source IPv4 addresses.

To configure the NAPT in IPv4 networks:

1. In configuration mode, go to the **[edit services]** hierarchy level.

```
[edit]
user@host# edit services
```

2. Configure the service set and NAT rule.

```
[edit services]
user@host# set service-set service-set-name nat-rules rule-name
```

In the following example, the name of the service set is **s1** and the name of the NAT rule is **rule-napt-44**.

```
[edit services]
user@host# set service-set s1 nat-rules rule-napt-44
```

3. Go to the **[interface-service]** hierarchy level of the service set.

```
[edit services]
user@host# edit service-set s1 interface-service
```

4. Configure the service interface.

```
[edit services service-set s1 interface service]
user@host# set service-interface service-interface-name
```

In the following example, the name of the service interface is **ms-0/1/0**.



NOTE: If the service interface is not present in the router, or the specified interface is not functional, the following command can result in an error.

```
[edit services service-set s1 interface service]
user@host# set service-interface ms-0/1/0
```

5. Go to the **[edit services nat]** hierarchy level. Issue the command from the top of the services hierarchy, or use the **top** keyword.

```
[edit services service-set s1 interface service]
user@host# top edit services nat
```

6. Configure the NAT pool with an address.

```
[edit services nat]
user@host# set pool pool-name address address
```

In the following example, the name of the pool is **napt-pool** and the address is **10.10.10.0**.

```
[edit services nat]
user@host# set pool napt-pool address 10.10.10.0
```

7. Configure the port.

```
[edit services nat]
user@host# set pool pool-name port port-type
```

In the following example, the port type is selected as **sequential** or **auto**.

```
[edit services nat]
user@host# set pool napt-pool port automatic
```



NOTE: Starting with Junos OS release 14.2, the **sequential** option is introduced to enable you to configure sequential allocation of ports. The **sequential** and **random-allocation** options available with the **port automatic** statement at the `[edit services nat pool nat-pool-name]` hierarchy level are mutually exclusive. You can include the **sequential** option for sequential allocation and the **random-allocation** option for random delegation of ports. By default, sequential allocation of ports takes place if you include only the **port automatic** statement at the `[edit services nat pool nat-pool-name]` hierarchy level. The **auto** option is hidden and is deprecated in Junos OS Release 14.2 and later, and is only maintained for backward compatibility. It might be removed completely in a future software release.

8. Configure the rule and the match direction.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction
```

In the following example, the name of the rule is **rule-napt-44** and the match direction is **input**.

```
[edit services nat]
user@host# set rule rule-napt-44 match-direction input
```

9. Configure the term, the action for the translated traffic, and the translation type.

```
[edit services nat]
user@host# set rule rule-name term term-name then translated translated-action
translation-type translation-type
```

In the following example, the name of the term is **t1**, the action for the translated traffic is **translated**, the name of the source pool is **napt-pool**, and the translation type is **napt-44**.

```
[edit services nat]
```

```
user@host# set rule rule-napt-44 match-direction input term t1 then translated
source-pool napt-pool translation-type napt-44
```

10. Go to the `[edit services adaptive-services-pics]` hierarchy level. In the command, the `top` keyword ensures that the command is run from the top of the hierarchy.

```
[edit services nat]
user@host# top edit services adaptive-services-pics
```

11. Configure the trace options.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag tracing parameter
```

In the following example, the tracing parameter is configured as `all`.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag all
```

12. Verify the configuration by using the `show` command at the `[edit services]` hierarchy level.

```
[edit services]
user@host# show
service-set s1 {
  nat-rules rule-napt-44;
  interface-service {
    service-interface ms-0/1/0;
  }
}
nat {
  pool napt-pool {
    address 10.10.10.0/32;
    port {
      automatic;
    }
  }
  rule rule-napt-44 {
    match-direction input;
    term t1 {
      then {
        translated {
          source-pool napt-pool;
          translation-type {
            napt-44;
          }
        }
      }
    }
  }
}
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```

The following example configures the translation type as `napt-44`.

```
[edit services]
user@host# show
service-set s1 {
```

```
    nat-rules rule-napt-44;
    interface-service {
        service-interface ms-0/1/0;
    }
}
nat {
    pool napt-pool {
        address 10.10.10.0/32;
        port {
            automatic auto;
        }
    }
    rule rule-napt-44 {
        match-direction input;
        term t1 {
            then {
                translated {
                    source-pool napt-pool;
                    translation-type {
                        napt-44;
                    }
                }
            }
        }
    }
}
adaptive-services-pics {
    traceoptions {
        flag all;
    }
}
```

Dynamic Address Translation to a Small Pool with Fallback to NAT

The following configuration shows dynamic address translation from a large prefix to a small pool, translating a /24 subnet to a pool of 10 addresses. When the addresses in the source pool (**src-pool**) are exhausted, NAT is provided by the NAPT overload pool (**pat-pool**).

```
[edit services nat]
pool src-pool {
    address-range low 192.16.2.1 high 192.16.2.10;
}
pool pat-pool {
    address-range low 192.16.2.11 high 192.16.2.12;
    port automatic auto;
    rule myrule {
        match-direction input;
        term myterm {
            from {
                source-address 10.150.1.0/24;
            }
            then {
                translated {
                    source-pool src-pool;
                    overload-pool pat-pool;
                    translation-type napt-44;
                }
            }
        }
    }
}
```

**Dynamic Address
Translation with Small
Pool**

}
The following configuration shows dynamic address translation from a large prefix to a small pool, translating a /24 subnet to a pool of 10 addresses. Sessions from the first 10 host sessions are assigned an address from the pool on a first-come, first-served basis, and any additional requests are rejected. Each host with an assigned NAT can participate in multiple sessions.

```
[edit services nat]
pool my-pool {
  address-range low 10.10.10.1 high 10.10.10.10;
}
rule src-nat {
  match-direction input;
  term t1 {
    from {
      source-address 192.168.1.0/24;
    }
    then {
      translated {
        translation-type dynamic-nat44;
        source-pool my-pool;
      }
    }
  }
}
```

Configuring Dynamic Source Address and Port Translation for IPv6 Networks

Network Address Port Translation (NAPT) is a method by which many network addresses and their TCP/UDP ports are translated into a single network address and its TCP/UDP ports. This translation can be configured in both IPv4 and IPv6 networks. This section describes the steps for configuring NAPT in IPv6 networks. For information about configuring NAPT in IPv4 networks, see [“Configuring Dynamic Source Address and Port Translation in IPv4 Networks” on page 127](#).

To configure NAPT, you must configure a rule at the **[edit services nat]** hierarchy level for dynamically translating the source IPv6 addresses.

To configure NAPT in IPv6 networks:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

```
[edit]
user@host# edit services nat
```

2. Define the pool of IPv6 source addresses that must be used for dynamic translation. For NAPT, also specify port numbers when configuring the source pool.

```
[edit services nat]
user@host# set pool pool name address IPv6 source addresses
user@host# set pool pool name port source ports
```

For example:

```
[edit services nat]
```

```
user@host# set pool IPV6-NAPT-Pool address 2002::1/96
user@host# set pool IPV6-NAPT-Pool port automatic sequential
```

3. Define a NAT rule for translating the source addresses. To do this, set the **match-direction** statement of the rule as **input**. In addition, define a term that uses **napt-66** as the translation type for translating the addresses of the pool defined in the previous step.

```
[edit services nat]
user@host# set rule rule name match-direction input
user@host# set rule rule name term term name then translated source-pool pool name
user@host# set rule rule name term term name then translated translation-type napt-66
```

For example:

```
[edit services nat]
user@host# set rule IPV6-NAPT-Rule match-direction input
user@host# set rule IPV6-NAPT-Rule term t1 then translated source-pool
  IPV6-NAPT-Pool
user@host# set rule IPV6-NAPT-Rule term t1 then translated translation-type napt-66
```

4. Enter the **up** command to navigate to the **[edit services]** hierarchy level.

```
[edit services nat]
user@host# up
```

5. Define a service set to specify the services interface that must be used, and reference the NAT rule implemented for NAPT translation.

```
[edit services]
user@host# set service-set service-set name interface-service service interface
  services interface
user@host# set service-set service-set name nat-rules rule name
```

For example:

```
[edit services]
user@host# set service-set IPV6-NAPT-ServiceSet interface-service service interface
  ms-0/1/0
user@host# set service-set IPV6-NAPT-ServiceSet nat-rules IPV6-NAPT-Rule
```

6. Define the trace options for the adaptive services PIC.

```
[edit services]
user@host# set adaptive-services-pics traceoptions flag tracing parameter
```

For example:

```
[edit services]
user@host# set adaptive-services-pics traceoptions flag all
```

The following example configures dynamic source (address and port) translation or NAPT for an IPv6 network.

```
[edit services]
user@host# show
  service-set IPV6-NAPT-ServiceSet {
    nat-rules IPV6-NAPT-Rule;
    interface-service {
      service-interface ms-0/1/0;
    }
  }
```

```
}
nat {
  pool IPV6-NAPT-Pool {
    address 2002::1/96;
    port automatic sequential;
  }
  rule IPV6-NAPT-Rule {
    match-direction input;
    term term1 {
      then {
        translated {
          source-pool IPV6-NAPT-Pool;
          translation-type {
            napt-66;
          }
        }
      }
    }
  }
}
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```

Example: Configuring NAT with Port Translation

This example shows how to configure NAT with port translation.

- [Requirements on page 133](#)
- [Overview on page 133](#)
- [Configuring NAT with Port Translation on page 133](#)

Requirements

This example uses the following hardware and software components:

- An MX Series 3D Universal Edge router with a Services DPC or an M Series Multiservice Edge router with a services PIC
- A domain name server (DNS)
- Junos OS Release 11.4 or higher

Overview

This example shows a complete CGN NAT44 configuration and advanced options.

Configuring NAT with Port Translation

Step-by-Step Procedure

To configure the service set:

1. Configure a service set.

```
user@host# edit services service-set ss2
```

2. Specify the NAT rule to be used.

```
[edit services service-set ss2]  
host# set nat-rules r1
```

3. Specify the interface service.

```
[edit services service-set ss2]  
host# set interface-service service-interface sp-5/0/0
```

Results user@host# show services service-sets sset2

```
nat-rules r1;  
interface-service {  
    service-interface sp-5/0/0;  
}
```

**Related
Documentation**

-

Example: NAPT Configuration for the MS-MPC

This example shows how to configure network address translation with port translation (NAPT) on an MX series router using a MultiServices Modular Port Concentrator (MS-MPC) as a services interface card.

- [Requirements on page 134](#)
- [Overview on page 134](#)
- [Configuration on page 134](#)

Requirements

This example uses the following hardware and software components:

- MX-series router
- MultiServices Modular Port Concentrator (MS-MPC)
- Junos OS Release 13.2R1 or higher

Overview

A service provider has chosen an MS-MPC as a platform to provide NAT services to accommodate new subscribers.

Configuration

To configure NAPT⁴⁴ using the MS-MPC as a services interface card, perform these tasks:

- [Configuring Interfaces on page 135](#)
- [Configure an Application Set of Acceptable ALG traffic on page 136](#)
- [Configuring a Stateful Firewall Rule on page 136](#)

- [Configuring NAT Pool and Rule on page 137](#)
- [Configuring the Service Set on page 138](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

```
set interfaces ge-0/2/0 unit 0 family inet address 10.255.248.2/24
set interfaces xe-1/1/0 unit 0 family inet address 10.255.247.2/24
set interfaces xe-1/1/0 unit 0 family inet service input service-set sset1
set interfaces xe-1/1/0 unit 0 family inet service output service-set sset1
set interfaces ms-3/0/0 unit 0 family inet
set applications application-set accept-algs application junos-http
set applications application-set accept-algs application junos-ftp
set applications application-set accept-algs application junos-tftp
set applications application-set accept-algs application junos-telnet
set applications application-set accept-algs application junos-sip
set applications application-set accept-algs application junos-rtcp
set services stateful-firewall rule sf-rule1 match-direction input-output
set services stateful-firewall rule sf-rule1 term sf-term1 from source-address
  10.255.247.0/24
set services stateful-firewall rule sf-rule1 term sf-term1 from application-sets accept-algs
set services stateful-firewall rule sf-rule1 term sf-term1 then accept
set services nat pool napt-pool address 1.1.1.0/24
set services nat pool napt-pool port automatic
* nat rule for napt
set services nat rule nat-rule1 match-direction input
set services nat rule nat-rule1 term nat-term1 from source-address 10.255.247.0/24
set services nat rule nat-rule1 term nat-term1 from application-sets accept-algs
set services nat rule nat-rule1 term nat-term1 then translated source-pool napt-pool
set services nat rule nat-rule1 term nat-term1 then translated translation-type napt-44
* nat rule for basic nat
set services service-set sset1 stateful-firewall-rules sf-rule1
set services service-set sset1 nat-rules nat-rule1
set services service-set sset1 interface-service service-interface ms-3/0/0
```

[Configuring Interfaces](#)

Step-by-Step Procedure

Configure the interfaces required for NAT processing. You will need the following interfaces:

- A customer-facing interface that handles traffic from and to the customer.
- An internet-facing interface.
- A services interface that provides NAT and stateful firewall services to the customer-facing interface

1. Configure the interface for the customer-facing interface.

```
user@host# edit
[edit ]
user@host# set interfaces xe-1/1/0 unit 0 family inet address 10.255.247.2/24
user@host# set interfaces xe-1/1/0 unit 0 family inet service input service-set sset1
user@host# set interfaces xe-1/1/0 unit 0 family inet service output service-set sset1
```

2. Configure the interface for the Internet-facing interface.

```
[edit ]  
set interfaces ge-0/2/0 unit 0 family inet address 10.255.248.2/24
```
3. Configure the interface for the service set that will connect services to the customer-facing interface. In our example, the interface resides on an MS-MPC.

```
[edit ]  
user@host# set interfaces ms-3/0/0 unit 0 family inet
```

Configure an Application Set of Acceptable ALG traffic

Step-by-Step Procedure

Identify the acceptable ALGs for incoming traffic.

1. Specify an application set that contains acceptable incoming ALG traffic.

```
user@host# set applications application-set accept-algs application junos-http  
user@host# set applications application-set accept-algs application junos-ftp  
user@host# set applications application-set accept-algs application junos-tftp  
user@host# set applications application-set accept-algs application junos-telnet  
user@host# set applications application-set accept-algs application junos-sip  
user@host# set applications application-set accept-algs application junos-rtcp
```

Results

```
user@host#edit services applications application-set accept-algs  
user@host#show  
application junos-http;  
application junos-ftp;  
application junos-tftp;  
application junos-telnet;  
application junos-sip;  
application junos-
```

Configuring a Stateful Firewall Rule

Step-by-Step Procedure

Configure a stateful firewall rule that will accept all incoming traffic.

1. Specify firewall matching for all input and output

```
user@host# set services stateful-firewall rule sf-rule1 match-direction input-output
```
2. Identify source-address and acceptable ALG traffic from the customer-facing interface.

```
user@host# set services stateful-firewall rule sf-rule1 term sf-term1 from  
source-address 10.255.247.0/24  
user@host# set services stateful-firewall rule sf-rule1 term sf-term1 from  
application-sets accept-algs  
user@host# set services stateful-firewall rule sf-rule1 term sf-term1 then accept
```

Results

```
user@host# edit services stateful-firewall
user@host# show
rule sf-rule1 {
  match-direction input-output;
  term sf-term1 {
    from {
      source-address {
        10.255.247.0/24;
      }
      application-sets accept-algs;
    }
    then {
      accept;
    }
  }
}
```

Configuring NAT Pool and Rule

Step-by-Step Procedure Configure a NAT pool and rule for address translation with automatic port assignment.

1. Configure the NAT pool with automatic port assignment.

```
user@host# set services nat pool napt-pool address 1.1.1.0/24
user@host# set services nat pool napt-pool port automatic auto
```

2. Configure a NAT rule that applies translation type **napt-44** using the defined NAT pool.

```
user@host# set services nat rule nat-rule1 term nat-term1 from application-sets
accept-algs
user@host# set services nat rule nat-rule1 term nat-term1 then translated source-pool
napt-pool
user@host# set services nat rule nat-rule1 term nat-term1 then translated
translation-type napt-44
```

Results

```
user@host#edit services nat
user@host#show

pool napt-pool {
    address 1.1.1.0/24;
    port {
        automatic;
    }
}
rule nat-rule1 {
    match-direction input;
    term nat-term1 {
        from {
            source-address {
                10.255.247.0/24;
            }
            application-sets accept-algs;
        }
        then {
            translated {
                source-pool napt-pool;
                translation-type {
                    napt-44;
                }
            }
        }
    }
}
```

Configuring the Service Set

Step-by-Step Procedure Configure an interface type service set.

1. Specify the NAT and stateful firewall rules that apply to customer traffic.

```
user@host set services service-set sset1 stateful-firewall-rules sf-rule1
user@host set services service-set sset1 nat-rules bat-rule1
```
2. Specify the services interface that applies the rules to customer traffic.

```
set services service-set sset1 interface-service service-interface ms-3/0/0
```

Results

```
user@host# edit services service-set sset1
user@host# show
set services service-set sset1 stateful-firewall-rules sf-rule1
set services service-set sset1 nat-rules nat-rule1
set services service-set sset1 interface-service service-interface ms-3/0/0
```

Related Documentation

- *Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards*

Example: Dynamic Source NAT as a Next-Hop Service

The following example shows dynamic-source NAT applied as a next-hop service:

```
[edit interfaces]
ge-0/2/0 {
```

```
    unit 0 {
        family mpls;
    }
}
sp-1/3/0 {
    unit 0 {
        family inet;
    }
    unit 20 {
        family inet;
    }
    unit 32 {
        family inet;
    }
}
[edit routing-instances]
protected-domain {
    interface ge-0/2/0.0;
    interface sp-1/3/0.20;
    instance-type vrf;
    route-distinguisher 10.58.255.17:37;
    vrf-import protected-domain-policy;
    vrf-export protected-domain-policy;
    routing-options {
        static {
            route 0.0.0.0/0 next-hop sp-1/3/0.20;
        }
    }
}
[edit policy-options]
policy-statement protected-domain-policy {
    term t1 {
        then reject;
    }
}
[edit services]
stateful-firewall {
    rule allow-all {
        match-direction input;
        term t1 {
            then {
                accept;
            }
        }
    }
}
nat {
    pool my-pool {
        address 10.58.16.100;
        port automatic;
    }
    rule hide-all {
        match-direction input;
        term t1 {
            then {
                translated {
```

```
        source-pool my-pool;
        translation-type napt-44;
    }
}
}
}
service-set null-sfw-with-nat {
    stateful-firewall-rules allow-all;
    nat-rules hide-all;
    next-hop-service {
        inside-service-interface sp-1/3/0.20;
        outside-service-interface sp-1/3/0.32;
    }
}
```

Securing Traffic Using NAT-PT and ALGs

- ALGs Available by Default for Junos OS Address Aware NAT on page 141
- Configuring Protocol Translation Between IPv6 and IPv4 Networks - NAT-PT on page 143
- Example: Configuring NAT-PT on page 150

ALGs Available by Default for Junos OS Address Aware NAT

The following application-level gateways (ALGs) listed in [Table 12 on page 141](#) are supported for NAT processing on the listed platforms.

To view the implementation details (port, protocol, and so on) for these Junos OS default applications, locate the Junos OS Default ALG Name in the table and then look up the listed name in the **groups**. For example, for details about TFTP, look up **junos-tftp** as shown.



TIP: The Junos OS provides the **junos-alg**, which enables other ALGs to function by handling ALG registrations, causing slow path packets to flow through registered ALGs, and transferring ALG events to the ALG plug-ins. The **junos-alg** ALG is automatically available on the MS-MPC and MS-MIC platforms and does not require further configuration.

```
user@host# show groups junos-defaults applications application junos-tftp
application-protocol tftp;
protocol udp;
destination-port 69;
```

Table 12: ALGs Available by Default

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
Basic TCP ALG	yes	yes	NOTE: Specific Junos ALGs are not supported. However, a feature called TCP tracker, available by default, performs segment ordering and retransmit and connection tracking, validations for TCP connections.
Basic UDP ALG	yes	yes	NOTE: TCP tracker performs limited integrity and validation checks for UDP.

Table 12: ALGs Available by Default (*continued*)

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
BOOTP	yes	no	<ul style="list-style-type: none"> • junos-bootpc • junos-bootps
DCE RPC Services	yes	yes	<ul style="list-style-type: none"> • junos-dce-rpc-portmap • junos-dcerpc-endpoint-mapper-service • junos-dcerpc-msexchange-directory-nsp • junos-dcerpc-msexchange-directory-rfr • junos-dcerpc-msexchange-information-store
DNS	yes	yes	<ul style="list-style-type: none"> • junos-dns-tcp • junos-dns-udp
FTP	yes	yes	<ul style="list-style-type: none"> • junos-ftp
H323	yes	no	<ul style="list-style-type: none"> • junos-h323
ICMP	yes	yes NOTE: ICMP messages are handled by default, but PING ALG support is not provided.	<ul style="list-style-type: none"> • junos-icmp-all • junos-icmp-ping
IIOp	yes	no	<ul style="list-style-type: none"> • junos-iiop-java • junos-iiop-orbix
IP	yes	The TCP tracker, available by default on these platforms, performs limited integrity and validation checks.	<ul style="list-style-type: none"> • junos-ip
NETBIOS	yes	no	<ul style="list-style-type: none"> • junos-netbios-datagram • junos-netbios-name-tcp • junos-netbios-name-udp • junos-netbios-session
NETSHOW	yes	no	<ul style="list-style-type: none"> • junos-netshow
PPTP	yes	yes	<ul style="list-style-type: none"> • junos-pptp
REALAUDIO	yes	no	<ul style="list-style-type: none"> • junos-realaudio
Sun RPC and RPC Port Map Services	yes	yes	<ul style="list-style-type: none"> • junos-rpc-portmap-tcp • junos-rpc-portmap-udp
RTSP	yes	yes	<ul style="list-style-type: none"> • junos-rtsp

Table 12: ALGs Available by Default (*continued*)

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
SIP	yes	Yes	<ul style="list-style-type: none"> • junos-sip <p>The SIP callid is <i>not</i> translated in register messages.</p> <p>NOTE: SIP sessions are limited to 12 hours (720 minutes) for NAT processing on the MS-MIC and MS-MPC interface cards. There is no time limit for SIP sessions on the MS-DPC.</p>
SNMP	yes	No	<ul style="list-style-type: none"> • junos-snmp-get • junos-snmp-get-next • junos-snmp-response junos-snmp-trap
SQLNET	yes	yes	<ul style="list-style-type: none"> • junos-sqlnet
TFTP	yes	yes	<ul style="list-style-type: none"> • junos-tftp
Traceroute	yes	no	<ul style="list-style-type: none"> • junos-traceroute
Unix Remote Shell Service	yes	Yes	<ul style="list-style-type: none"> • junos-rsh
WINFrame	yes	No	<ul style="list-style-type: none"> • junos-citrix-winframe • junos-citrix-winframe-udp
TALK-UDP	No	Yes	<ul style="list-style-type: none"> • junos-talk-udp
MS RPC	No	Yes	<ul style="list-style-type: none"> • junos-rpc-portmap-tcp • junos-rpc-portmap-udp • junos-rpc-services-tcp • junos-rpc-services-udp

Related Documentation • [ALG Descriptions on page 277](#)

Configuring Protocol Translation Between IPv6 and IPv4 Networks - NAT-PT

To configure the translation type as **basic-nat-pt**, you must configure the DNS ALG application, the NAT pools and rules, a service set with a service interface, and trace options. This topic includes the following tasks:

- [Configuring the DNS ALG Application on page 144](#)
- [Configuring the NAT Pool and NAT Rule on page 144](#)
- [Configuring the Service Set for NAT on page 147](#)
- [Configuring Trace Options on page 148](#)

Configuring the DNS ALG Application

To configure the DNS ALG application:

1. In configuration mode, go to the **[edit applications]** hierarchy level.

```
[edit]
user@host# edit applications
```

2. Configure the ALG to which the DNS traffic is destined at the **[edit applications]** hierarchy level. Define the application name and specify the application protocol to use in match conditions in the first NAT rule or term.

```
[edit applications]
user@host# set application application-name application-protocol application-protocol
```

In the following example, the application name is **dns-alg** and application protocol is **dns**.

```
[edit applications]
user@host# set application dns-alg application-protocol dns
```

3. Verify the configuration by using the **show** command at the **[edit applications]** hierarchy level.

```
[edit applications]
user@host# show
application dns-alg {
    application-protocol dns;
}
```

Configuring the NAT Pool and NAT Rule

To configure the NAT pool and NAT rule:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

```
[edit]
user@host# edit services nat
```

2. Configure the NAT pool and its address.

```
[edit services nat]
user@host# set pool pool-name address address
```

In the following example, the name of the NAT pool is **p1** and the address is **10.10.10.2/32**.

```
[edit services nat]
user@host# set pool p1 address 10.10.10.2/32
```

3. Configure the source pool and its address.

```
[edit services nat]
user@host# set pool source-pool-name address address
```

In the following example, the name of the source pool is **src_pool0** and the source pool address is **20.1.1.1/32**.

```
[edit services nat]
```

```
user@host# set pool src_pool0 address 20.1.1.1/32
```

4. Configure the destination pool and its address.

```
[edit services nat]
user@host# set pool destination-pool-name address address
```

In the following example, the name of the destination pool is **dst_pool0** and the destination pool address is **50.1.1.2/32**.

```
[edit services nat]
user@host# set pool dst_pool0 address 50.1.1.2/32
```

5. Configure the rule and the match direction.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction
```

In the following example, the rule name is **rule-basic-nat-pt** and the match direction is **input**.

```
[edit services nat]
user@host# set rule basic-nat-pt match-direction input
```

6. Configure the term and the input conditions for the NAT term.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term term from from
```

In the following example, the term is **t1** and the input conditions are **source-address 2000::2/128**, **destination-address 4000::2/128**, and **applications dns_alg**.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 from source-address 2000::2/128
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 from destination-address 4000::2/128
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 from applications dns_alg
```

7. Configure the NAT term action and the properties of the translated traffic.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 then term-action translated-property
```

In the following example, the term action is **translated** and the properties of the translated traffic are **source-pool src_pool0**, **destination-pool dst_pool0**, and **dns-alg-prefix 10:10:10::0/96**.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 then translated source-pool src_pool0
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 then translated destination-pool
dst_pool0
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 then translated dns-alg-prefix
10:10:10::0/96
```

8. Configure the translation type.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 then translated translation-type
translation-type
```

In the following example, the translation type is **basic-nat-pt**.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 then translated translation-type
basic-nat-pt
```

9. Configure another term and the input conditions for the NAT term.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term term-name from from
```

In the following example, the term name is **t2** and the input conditions are **source-address 2000::2/128** and **destination-address 10:10:10::0/96**.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t2 from source-address 2000::2/128
[edit services nat]
user@host# set rule rule-basic-nat-pt term t2 from destination-address 10:10:10::0/96
```

10. Configure the NAT term action and the property of the translated traffic.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t2 then term-action translated-property
```

In the following example, the term action is **translated** and the property of the translated traffic is **source-prefix 19.19.19.1/32**.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t2 then translated source-prefix
19.19.19.1/32
```

11. Configure the translation type.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t2 then translated translation-type
translation-type
```

In the following example, the translation type is **basic-nat-pt**.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t2 then translated translation-type
basic-nat-pt
```

12. Verify the configuration by using the **show** command at the **[edit services nat]** hierarchy level.

```
[edit services nat]
user@host# show
pool p1 {
    address 10.10.10.2/32;
}
pool src_pool0 {
    address 20.1.1.1/32;
}
pool dst_pool0 {
    address 50.1.1.2/32;
}
rule rule-basic-nat-pt {
    match-direction input;
    term t1 {
        from {
            source-address {
```

```

        2000::2/128;
    }
    destination-address {
        4000::2/128;
    }
    applications dns_alg;
}
then {
    translated {
        source-pool src_pool0;
        destination-pool dst_pool0;
        dns-alg-prefix 10:10:10::0/96;
        translation-type {
            basic-nat-pt;
        }
    }
}
}
term t2 {
    from {
        source-address {
            2000::2/128;
        }
        destination-address {
            10:10:10::0/96;
        }
    }
    then {
        translated {
            source-prefix 19.19.19.1/32;
            translation-type {
                basic-nat-pt;
            }
        }
    }
}
}
}

```

Configuring the Service Set for NAT

To configure the service set for NAT:

1. In configuration mode, go to the **[edit services]** hierarchy level.

```

[edit]
user@host# edit services

```

2. Configure the service set.

```

[edit services]
user@host# edit service-set service-set-name

```

In the following example, the name of the service set is **ss_dns**.

```

[edit services]
user@host# edit service-set ss_dns

```

3. Configure the service set with NAT rules.

```

[edit services service-set ss_dns]
user@host# set nat-rules rule-name

```

In the following example, the rule name is **rule-basic-nat-pt**.

```
[edit services service-set ss_dns]
user@host# set nat-rules rule-basic-nat-pt
```

4. Configure the service interface.

```
[edit services service-set ss_dns]
user@host# set interface-service service-interface service-interface-name
```

In the following example, the name of service interface is **sp-1/2/0**.

```
[edit services service-set ss_dns]
user@host# set interface-service service-interface sp-1/2/0
```

5. Verify the configuration by using the **show services** command from the **[edit]** hierarchy level.

```
[edit]
user@host# show services
  service-set ss_dns {
    nat-rules rule-basic-nat-pt;
    interface-service {
      service-interface sp-1/2/0;
    }
  }
```

Configuring Trace Options

To configure the trace options:

1. In configuration mode, go to the **[edit services adaptive-services-pics]** hierarchy level.

```
[edit]
user@host# edit services adaptive-services-pics
```

2. Configure the trace options.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag tracing parameter
```

In the following example, the tracing parameter is **all**.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag all
```

3. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```

The following example configures the translation type as **basic-nat-pt**.

```
[edit]
user@host# show services
```

```

service-set ss_dns {
    nat-rules rule-basic-nat-pt;
    interface-service {
        service-interface sp-1/2/0;
    }
}
nat {
    pool p1 {
        address 10.10.10.2/32;
    }
    pool src_pool0 {
        address 20.1.1.1/32;
    }
    pool dst_pool0 {
        address 50.1.1.2/32;
    }
    rule rule-basic-nat-pt {
        match-direction input;
        term t1 {
            from {
                source-address {
                    2000::2/128;
                }
                destination-address {
                    4000::2/128;
                }
                applications dns_alg;
            }
            then {
                translated {
                    source-pool src_pool0;
                    destination-pool dst_pool0;
                    dns_alg-prefix 10:10:10::0/96;
                    translation-type {
                        basic-nat-pt;
                    }
                }
            }
        }
        term t2 {
            from {
                source-address {
                    2000::2/128;
                }
                destination-address {
                    10:10:10::0/96;
                }
            }
            then {
                translated {
                    source-prefix 19.19.19.1/32;
                    translation-type {
                        basic-nat-pt;
                    }
                }
            }
        }
    }
}
adaptive-services-pics {
    traceoptions {

```

```
        flag all;  
    }  
}
```

Example: Configuring NAT-PT

A Domain Name System application-level gateway (DNS ALG) is used with Network Address Translation-Protocol Translation (NAT-PT) to facilitate name-to-address mapping. You can configure the DNS ALG to map addresses returned in the DNS response to an IPv6 address.

When you configure NAT-PT with DNS ALG support, you must configure two NAT rules or one rule with two terms. In this example, you configure two rules. The first NAT rule ensures that the DNS query and response packets are translated correctly. For this rule to work, you must configure a DNS ALG application and reference it in the rule. The second rule is required to ensure that NAT sessions are destined to the address mapped by the DNS ALG.

Then, you must configure a service set, and then apply the service set to the interfaces.

This example describes how to configure NAT-PAT with DNS ALG:

- [Requirements on page 150](#)
- [Overview and Topology on page 150](#)
- [Configuration of NAT-PT with DNS ALGs on page 152](#)

Requirements

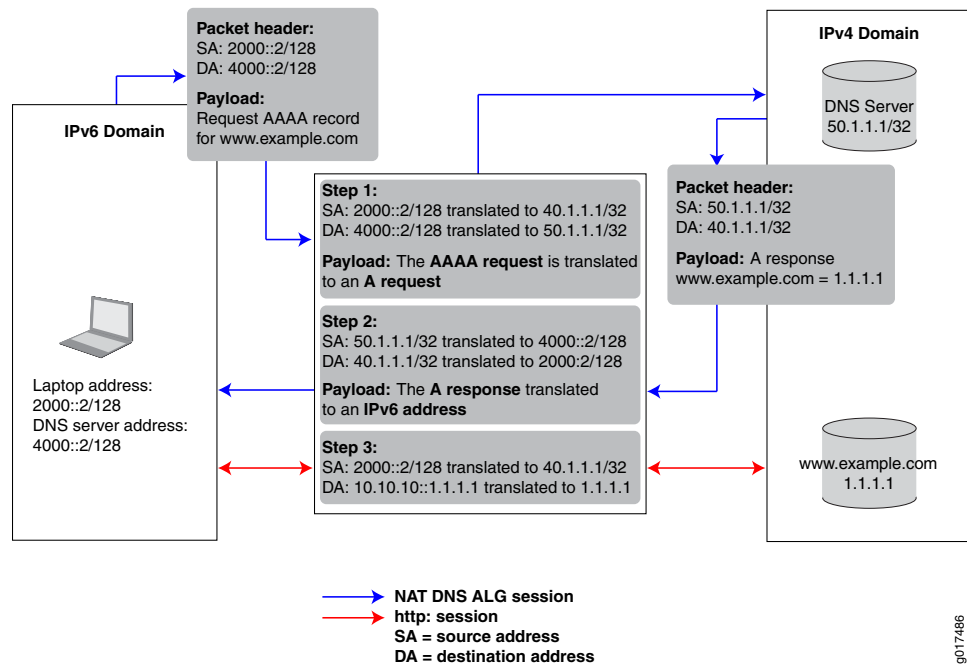
This example uses the following hardware and software components:

- Junos OS Release 11.2
- A multiservices interface (**ms-**)

Overview and Topology

The following scenario shows the process of NAT-PT with DNS ALG when a laptop in an IPv6-only domain requests access to a server in an IPv4-only domain.

Figure 8: Configuring DNS ALGs with NAT-PT Network Topology



The Juniper Networks router in the center of the illustration performs address translation in two steps. When the laptop requests a session with the **www.example.com** server that is in an IPv4-only domain, the Juniper Networks router performs the following:

- Translates the IPv6 laptop and DNS server addresses into IPv4 addresses.
- Translates the AAAA request from the laptop into an A request so that the DNS server can provide the IPv4 address.

When the DNS server responds with the A request, the Juniper Networks router performs the following:

- Translates the IPv4 DNS server address back into an IPv6 address.
- Translates the A request back into a AAAA request so that the laptop now has the 96-bit IPv6 address of the **www.example.com** server.

After the laptop receives the IPv6 version of the **www.example.com** server address, the laptop initiates a second session using the 96-bit IPv6 address to access that server. The Juniper Networks router performs the following:

- Translates the laptop IPv4 address directly into its IPv4 address.
- Translates the 96-bit IPv6 **www.example.com** server address into its IPv4 address.

Configuration of NAT-PT with DNS ALGs

To configure NAT-PT with DNS ALG, perform the following tasks:

- [Configuring the Application-Level Gateway on page 152](#)
- [Configuring the NAT Pools on page 153](#)
- [Configuring the DNS Server Session: First NAT Rule on page 154](#)
- [Configuring the HTTP Session: Second NAT Rule on page 157](#)
- [Configuring the Service Set on page 159](#)
- [Configuring the Stateful Firewall Rule on page 161](#)
- [Configuring Interfaces on page 162](#)

Configuring the Application-Level Gateway

Step-by-Step Procedure

Configure the DNS application as the ALG to which the DNS traffic is destined. The DNS application protocol closes the DNS flow as soon as the DNS response is received. When you configure the DNS application protocol, you must specify the UDP protocol as the network protocol to match in the application definition.

To configure the DNS application:

1. In configuration mode, go to the **[edit applications]** hierarchy level.

```
user@host# edit applications
```

2. Define the application name and specify the application protocol to use in match conditions in the first NAT rule.

```
[edit applications]
user@host# set application application-name application-protocol protocol-name
```

For example:

```
[edit applications]
user@host# set application dns_alg application-protocol dns
```

3. Specify the protocol to match, in this case UDP.

```
[edit applications]
user@host# set application application-name protocol type
```

For example:

```
[edit applications]
user@host# set application dns_alg protocol udp
```

4. Define the UDP destination port for additional packet matching, in this case the domain port.

```
[edit applications]
user@host# set application application-name destination-port value
```

For example:

```
[edit applications]
user@host# set application dns_alg destination-port 53
```

Results [edit applications]
 user@host# show
 application dns_alg {
 application-protocol dns;
 protocol udp;
 destination-port 53;
 }

Configuring the NAT Pools

Step-by-Step Procedure In this configuration, you configure two pools that define the addresses (or prefixes) used for NAT. These pools define the IPv4 addresses that are translated into IPv6 addresses. The first pool includes the IPv4 address of the source. The second pool defines the IPv4 address of the DNS server. To configure NAT pools:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.
 user@host# edit services nat
2. Specify the name of the first pool and the IPv4 source address (laptop).
 [edit services nat]
 user@host# set pool *nat-pool-name* address *ip-prefix*

For example:

```
[edit services nat]
user@host# set pool pool1 address 40.1.1.1/32
```

3. Specify the name of the second pool and the IPv4 address of the DNS server.
 [edit services nat]
 user@host# set pool *nat-pool-name* address *ip-prefix*

For example:

```
[edit services nat]
user@host# set pool pool2 address 50.1.1.1/32
```

Results The following sample output shows the configuration of NAT pools.

```
[edit services nat]
user@host# show
pool pool1 {
  address 40.1.1.1/32;
}
pool pool2 {
  address 50.1.1.1/32;
}
```

Configuring the DNS Server Session: First NAT Rule

Step-by-Step Procedure The first NAT rule is applied to DNS traffic going to the DNS server. This rule ensures that the DNS query and response packets are translated correctly. For this rule to work, you must configure a DNS ALG application and reference it in the rule. The DNS application was configured in [“Configuring the DNS ALG Application” on page 144](#). In addition, you must specify the direction in which traffic is matched, the source address of the laptop, the destination address of the DNS server, and the actions to take when the match conditions are met.

To configure the first NAT rule:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

```
user@host# edit services nat
```

2. Specify the name of the NAT rule.

```
[edit services nat]
user@host# edit rule rule-name
```

For example:

```
[edit services nat]
user@host# edit rule rule1
```

3. Specify the name of the NAT term.

```
[edit services nat rule rule-name]
user@host# edit term term-name
```

For example:

```
[edit services nat rule rule1]
user@host# edit term term1
```

4. Define the match conditions for this rule.

- a. Specify the IPv6 source address of the device (laptop) attempting to access an IPv4 address.

```
[edit services nat rule rule-name term term-name]
user@host# set from source-address source-address
```

For example:

```
[edit services nat rule rule1 term term1]
user@host# set from source-address 2000::2/128
```

- b. Specify the IPv6 destination address of the DNS server.

```
[edit services nat rule rule-name term term-name]
user@host# set from destination-address prefix
```

For example:

```
[edit services nat rule rule1 term term1]
user@host# set from destination-address 4000::2/128
```

- c. Reference the DNS application to which the DNS traffic destined for port 53 is applied.

```
[edit services nat rule rule1 term term1]
user@host# set from applications application-name
```

In this example, the application name configured in the *Configuring the DNS Application* step is `dns_alg`:

```
[edit services nat rule rule1 term term1]
user@host# set from applications dns_alg
```

5. Define the actions to take when the match conditions are met. The source and destination pools you configured in “[Configuring the NAT Pools](#)” on page 153 are applied here.

- a. Apply the NAT pool configured for source translation.

```
[edit services nat rule rule-name term term-name]
user@host# set then translated source-pool nat-pool-name
```

For example:

```
[edit services nat rule rule1 term term1]
user@host# set then translated source-pool pool1
```

- b. Apply the NAT pool configured for destination translation.

```
[edit services nat rule rule-name term term-name]
user@host# set then translated destination-pool nat-pool-name
```

For example:

```
[edit services nat rule rule1 term term1]
user@host# set then translated source-pool pool2
```

6. Define the DNS ALG 96-bit prefix for IPv4-to-IPv6 address mapping.

```
[edit services nat rule rule-name term term-name]
user@host# set then translated dns-alg-prefix dns-alg-prefix
```

For example:

```
[edit services nat rule rule1 term term1]
user@host# set then translated dns-alg-prefix 10:10:10::0/96
```

7. Specify the type of NAT used for source and destination traffic.

```
[edit services nat rule rule-name term term-name]
user@host# set then translated translation-type basic-nat-pt
```

For example:

```
[edit services nat rule rule1 term term1]
user@host# set then translated translation-type basic-nat-pt
```



NOTE: In this example, since NAT is achieved using address-only translation, the `basic-nat-pt` translation type is used. To achieve NAT using address and port translation (NAPT), use the `napt-pt` translation type.

8. Specify the direction in which to match traffic that meets the rule conditions.

```
[edit services nat rule rule-name]  
user@host# set match-direction (input | output)
```

For example:

```
[edit services nat rule rule1]  
user@host# set match-direction input
```

9. Configure system logging to record information from the services interface to the `/var/log` directory.

```
[edit services nat rule rule-name term term-name]  
user@host# set then syslog
```

For example:

```
[edit services nat rule rule1 term term1]  
user@host# set then syslog
```

Results The following sample output shows the configuration of the first NAT rule that goes to the DNS server.

```
[edit services nat]
user@host# show
rule rule1 {
  match-direction input;
  term term1 {
    from {
      source-address {
        2000::2/128;
      }
      destination-address {
        4000::2/128;
      }
      applications dns_alg;
    }
    then {
      translated {
        source-pool pool1;
        destination-pool pool2;
        dns-alg-prefix 10:10:10::0/96;
        translation-type {
          basic-nat-pt;
        }
      }
    }
    syslog;
  }
}
```

Configuring the HTTP Session: Second NAT Rule

Step-by-Step Procedure

The second NAT rule is applied to destination traffic going to the IPv4 server (www.example.com). This rule ensures that NAT sessions are destined to the address mapped by the DNS ALG. For this rule to work, you must configure the DNS ALG address map that correlates the DNS query or response processing done by the first rule with the actual data sessions processed by the second rule. In addition, you must specify the direction in which traffic is matched: the IPv4 address for the IPv6 source address (laptop), the 96-bit prefix to prepend to the IPv4 destination address (www.example.com), and the translation type.

To configure the second NAT rule:

1. In configuration mode, go to the following hierarchy level.

```
user@host# edit services nat
```

2. Specify the name of the NAT rule and term.

```
[edit services nat]
user@host# edit rule rule-name term term-name
```

For example:

```
[edit services nat]
user@host# edit rule rule2 term term1
```

3. Define the match conditions for this rule:
 - a. Specify the IPv6 address of the device attempting to access the IPv4 server.

```
[edit services nat rule rule-name term term-name]  
user@host# set from source-address source-address
```

For example:

```
[edit services nat rule rule2 term term1]  
user@host# set from source-address 2000::2/128
```

- b. Specify the 96-bit IPv6 prefix to prepend to the IPv4 server address.

```
[edit services nat rule rule-name term term-name]  
user@host# set from destination-address prefix
```

For example:

```
[edit services nat rule rule2 term term1]  
user@host# set from destination-address 10:10:10::c0a8:108/128
```

4. Define the actions to take when the match conditions are met.
 - Specify the prefix for the translation of the IPv6 source address.

```
[edit services nat rule rule-name term term-name]  
user@host# set then translated source-prefix source-prefix
```

For example:

```
[edit services nat rule rule2 term term1]  
user@host# set then translated source-prefix 19.19.19.1/32
```

5. Specify the type of NAT used for source and destination traffic.

```
[edit services nat rule rule-name term term-name]  
user@host# set then translated translation-type basic-nat-pt
```

For example:

```
[edit services nat rule rule2 term term1]  
user@host# set then translated translation-type basic-nat-pt
```



NOTE: In this example, since NAT is achieved using address-only translation, the *basic-nat-pt* translation type is used. To achieve NAT using address and port translation (NAPT), you must use the *napt-pt* translation type.

6. Specify the direction in which to match traffic that meets the conditions in the rule.

```
[edit services nat rule rule-name]  
user@host# set match-direction (input | output)
```

For example:

```
[edit services nat rule rule2]  
user@host# set match-direction input
```


Results The following sample output shows the configuration of the second NAT rule.

```
[edit services nat]
user@host# show
rule rule2 {
  match-direction input;
  term term1 {
    from {
      source-address {
        2000::2/128;
      }
      destination-address {
        10:10:10::c0a8:108/128;
      }
    }
    then {
      translated {
        source-prefix 19.19.19.1/32;
        translation-type {
          basic-nat-pt;
        }
      }
    }
  }
}
```

Configuring the Service Set

Step-by-Step Procedure This service set is an interface service set used as an action modifier across the entire services (**ms-**) interface. Stateful firewall and NAT rule sets are applied to traffic processed by the services interface.

To configure the service set:

1. In configuration mode, go to the **[edit services]** hierarchy level.

```
user@host# edit services
```

2. Define a service set.

```
[edit services]
user@host# edit service-set service-set-name
```

For example:

```
[edit services]
user@host# edit service-set ss
```

3. Specify properties that control how system log messages are generated for the service set.

```
[edit services service-set ss]
user@host# set syslog host local services severity-level
```

The example below includes all severity levels.

```
[edit services service-set ss]
user@host# set syslog host local services any
```

4. Specify the stateful firewall rule included in this service set.

```
[edit services service-set ss]
user@host# set stateful-firewall-rules rule1 severity-level
```

The example below references the stateful firewall rule defined in [“Configuring the Stateful Firewall Rule” on page 161](#).

```
[edit services service-set ss]
user@host# set stateful-firewall-rules rule1
```

5. Define the NAT rules included in this service set.

```
[edit services service-set ss]
user@host# set nat-rules rule-name
```

The example below references the two rules defined in this configuration example.

```
[edit services service-set ss]
user@host# set nat-rules rule1
user@host# set nat-rules rule2
```

6. Configure an adaptive services interface on which the service is to be performed.

```
[edit services service-set ss]
user@host# set interface-service service-interface interface-name
```

For example:

```
[edit services service-set ss]
user@host# interface-service service-interface ms-2/0/0
```

Only the device name is needed, because the router software manages logical unit numbers automatically. The services interface must be an adaptive services interface for which you have configured **unit 0 family inet** at the `[edit interfaces interface-name]` hierarchy level in [“Configuring Interfaces” on page 162](#).

Results The following sample output shows the configuration of the service set.

```
[edit services]
user@host# show
service-set ss {
  syslog {
    host local {
      services any;
    }
  }
  stateful-firewall-rules rule1;
  nat-rules rule1;
  nat-rules rule2;
  interface-service {
    service-interface ms-2/0/0;
  }
}
```

Configuring the Stateful Firewall Rule

Step-by-Step Procedure This example uses a stateful firewall to inspect packets for state information derived from past communications and other applications. The NAT-PT router checks the traffic flow matching the direction specified by the rule, in this case both input and output. When a packet is sent to the services (**ms-**) interface, direction information is carried along with it.

To configure the stateful firewall rule:

1. In configuration mode, go to the **[edit services stateful firewall]** hierarchy level.

```
user@host# edit services stateful firewall
```

2. Specify the name of the stateful firewall rule.

```
[edit services stateful-firewall]
user@host# edit rule rule-name
```

For example:

```
[edit services stateful-firewall]
user@host# edit rule rule1
```

3. Specify the direction in which traffic is to be matched.

```
[edit services stateful-firewall rule rule-name]
user@host# set match-direction (input | input-output | output)
```

For example:

```
[edit services stateful-firewall rule rule1]
user@host# set match-direction input-output
```

4. Specify the name of the stateful firewall term.

```
[edit services stateful-firewall rule rule-name]
user@host# edit term term-name
```

For example:

```
[edit services stateful-firewall rule rule1]
user@host# edit term term1
```

5. Define the terms that make up this rule.

```
[edit services stateful-firewall rule rule-name term term-name]  
user@host# set then accept
```

For example:

```
[edit services stateful-firewall rule rule1 term term1]  
user@host# set then accept
```

Results The following sample output shows the configuration of the services stateful firewall.

```
[edit services]  
user@host# show  
stateful-firewall {  
  rule rule1 {  
    match-direction input-output;  
    term term1 {  
      then {  
        accept;  
      }  
    }  
  }  
}
```

Configuring Interfaces

Step-by-Step Procedure After you have defined the service set, you must apply services to one or more interfaces installed on the router. In this example, you configure one interface on which you apply the service set for input and output traffic. When you apply the service set to an interface, it automatically ensures that packets are directed to the services (**ms-**) interface.

To configure the interfaces:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level.

```
user@host# edit interfaces
```

2. Configure the interface on which the service set is applied to automatically ensure that packets are directed to the services (**ms-**) interface.

- a. For IPv4 traffic, specify the IPv4 address.

```
[edit interfaces]  
user@host# set ge-1/0/9 unit 0 family inet address 30.1.1.1/24
```

- b. Apply the service set defined in [“Configuring Interfaces” on page 162](#).

```
[edit interfaces]  
user@host# set ge-1/0/9 unit 0 family inet6 service input service-set ss  
user@host# set ge-1/0/9 unit 0 family inet6 service output service-set ss
```

- c. For IPv6 traffic, specify the IPv6 address.

```
[edit interfaces]  
user@host# set ge-1/0/9 unit 0 family inet6 address 2000::1/64
```

3. Specify the interface properties for the services interface that performs the service.

```
[edit interfaces]
user@host# set ms-2/0/0 services-options syslog host local services any
user@host# set ms-2/0/0 unit 0 family inet
user@host# set ms-2/0/0 unit 0 family inet6
```

Results The following sample output shows the configuration of the interfaces for this example.

```
[edit interfaces]
user@host# show

ge-1/0/9 {
  unit 0 {
    family inet {
      address 30.1.1.1/24;
    }
    family inet6 {
      service {
        input {
          service-set ss;
        }
        output {
          service-set ss;
        }
      }
      address 2000::1/64;
    }
  }
}

ms-2/0/0 {
  services-options {
    syslog {
      host local {
        services any;
      }
    }
  }
  unit 0 {
    family inet;
    family inet6;
  }
}
```

- Related Documentation**
- [Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards](#)
 - [Configuring Service Sets to be Applied to Services Interfaces on page 31](#)
 - [Example: Configuring Layer 3 Services and the Services SDK on Two PICs on page 339](#)
 - [dns-alg-prefix on page 1350](#)
 - [dns-alg-pool on page 1349](#)

Reducing Traffic and Bandwidth Requirements Using Port Control Protocol

- [Port Control Protocol Overview on page 165](#)
- [Configuring Port Control Protocol on page 167](#)
- [Example: Configuring Port Control Protocol with NAPT44 on page 170](#)

Port Control Protocol Overview

The Port Control Protocol (PCP) provides a mechanism to control the forwarding of incoming packets by upstream devices such as NAT44, and firewall devices, and a mechanism to reduce application keep-alive traffic. PCP is designed to be implemented in the context of both Carrier-Grade NATs (CGNs) and small NATs (for example, residential NATs). PCP allows hosts to operate servers for a long time (as in the case of a webcam) or a short time (for example, while playing a game or on a phone call) when behind a NAT device, including when behind a CGN operated by their Internet service provider. PCP allows applications to create mappings from an external IP address and port to an internal IP address and port. These mappings are required for successful inbound communications destined to machines located behind a NAT or a firewall. After creating a mapping for incoming connections, it is necessary to inform remote computers about the IP address and port for the incoming connection. This is usually done in an application-specific manner.

PCP consists of the following components:


- **PCP client**—A host or gateway that issues PCP requests to a PCP server in order to obtain and control resources.
- **PCP server**—Typically a CGN gateway or co-located server that receives and processes PCP requests

Many NAT-friendly applications send frequent application-level messages to ensure their session are not be timed out by a NAT. These applications can reduce the frequency of such NAT keep-alive messages by using PCP to learn and influence the NAT mapping lifetime. This helps reduce bandwidth on the subscriber's access network, traffic to the server, and battery consumption on mobile devices.

The Junos OS enables configuring PCP servers for mapping flows using NAPT44 capabilities such as port forwarding and port block allocation. Flows can be processed from these sources:


- Traffic containing PCP requests received directly from UEs as shown in [Figure 9 on page 166](#).

Figure 9: Basic PCP NAPT44 Topology



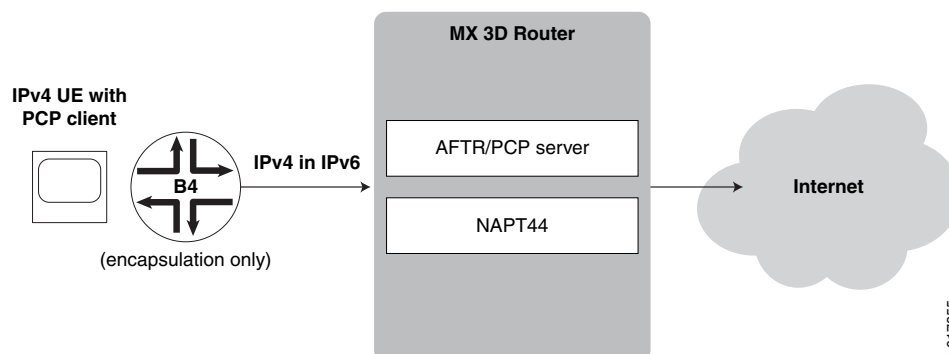
- Mapping of traffic containing PCP requests added by a router functioning as a DS-Lite software initiator (B4). This mode, known as *DS-Lite plain mode*, is shown in [Figure 10 on page 166](#)

Figure 10: PCP with DS-Lite Plain Mode



- Mapping of traffic containing PCP requests initiated directly by UEs and encapsulated by a router functioning as a DS-Lite software initiator (B4). This mode, known as *DS-Lite tunnel mode*, is shown in [Figure 11 on page 167](#).

Figure 11: PCP with DS-Lite Tunnel Mode



NOTE: The Junos OS does not support deterministic port block allocation for PCP-originated traffic.

Related Documentation

- [Configuring Port Control Protocol on page 167](#)

Configuring Port Control Protocol

This topic describes the following configuration tasks:

- [Configuring PCP Server Options on page 167](#)
- [Configuring a PCP Rule on page 168](#)
- [Configuring a Service Set to Apply PCP on page 169](#)
- [SYSLOG Message Configuration on page 169](#)

Configuring PCP Server Options

1. Go to the `[edit services pcp pcp-server server-name]` hierarchy level and specify a PCP server name.

```
user @host# edit services pcp pcp-server server-name
```

2. Set the IPv4 or IPv6 addresses of the server. For PCP DS-Lite, the **ipv6-address** must match the address of the AFTR (Address Family Transition Router or software concentrator).

```
[edit services pcp pcp-server s1]
user @host# set ipv6-address ipv6-address
```

or

```
[edit services pcp pcp-server s1]
user @host# set ipv4-address ipv4-address
```

3. For PCP DS-Lite, provide the name of the DS-Lite software concentrator configuration.

```
[edit services pcp pcp-server s1]
user @host# set software-concentrator software-concentrator-name
```

- Specify the minimum and maximum mapping lifetimes for the server.

```
[edit services pcp pcp-server s1]
user @host# set mapping-lifetime-minimum mapping-lifetime-minimum
user @host# set mapping-lifetime-maximum mapping-lifetime-maximum
```

- Specify the time limits for generating short lifetime or long lifetime errors.

```
[edit services pcp pcp-server s1]
user @host# set short-lifetime-error short-lifetime-error
user @host# set long-lifetime-error long-lifetime-error
```

- (Optional)—Enable PCP options on the specified PCP server. The following options are available—**third-party** and **prefer-failure**. The third-party option is required to enable third-party requests by the PCP client. DS-Lite requires the **third-party** option. The **prefer-failure** option requests generation of an error message when the PCP client requests a specific IP address/port that is not available, rather than assigning another available address from the NAT pool. If **prefer-failure** is not specified NAPT44 assigns an available address/port from the NAT pool based on the configured NAT options.

```
[edit services pcp pcp-server s1]
user @host# set pcp-options third-party
user @host# set pcp-options prefer-failure
```

- (Optional)—Specify which NAT pool to use for mapping.

```
[edit services pcp pcp-server s1]
user @host# set nat-options pcp-nat-pool pool-name1 <poolname2...>
```



NOTE: When you do not explicitly specify a NAT pool for mapping, the Junos OS performs a partial rule match based on source IP, source port and protocol; the Junos OS uses the NAT pool configured for the first matching rule to allocate mappings for PCP.

You *must* use explicit configuration in order to use multiple NAT pools.

- (Optional)—Configure the maximum number of mappings per client. The default is 32 and maximum is 128.

```
[edit services pcp pcp-server s1]
user @host# set max-mappings-per-client max-mappings-per-client
```

Configuring a PCP Rule

A PCP rule is has the same basic options as all service set rules:

- A **term** option that allows a single rule to have multiple applications.
 - A **from** option that identifies the traffic that is subject to the rule.
 - A **then** option that identifies what action is to be taken. In the case of a PCP rule, this option identifies the pcp server that handles selected traffic
- Go to the **[edit services pcp rule *rulename*]** hierarchy level and specify **match-direction** input.

```
user @host# edit services pcp rule rulename
user @host# set match-direction input
```

2. Go to the `[edit services pcp rule rulename term termname]` hierarchy level and provide a *termname*.

```
user @host# edit term termname
```

3. (Optional)—Provide a **from** option to filter the traffic to be selected for processing by the rule. When you omit the **from** option, all traffic handled by the service set's service interface is subject to the rule.
4. Set the **then** option to identify the target pcp server.

```
[edit services pcp rule rulename term termname]
user @host# set then pcp-server server-name
```

Configuring a Service Set to Apply PCP

To use PCP, you must provide the rule-name (or name of a list of rulenames) in the **pcp-rule *rulename*** option.

1. Go to the `[edit services service-set service-set-name]` hierarchy level.

```
user @host# edit services service-set service-set-name
```

2. If this is a new service set, provide basic service set information, including interface information and any other rules that may apply.
3. Specify the name of the PCP rule or rule list used to send traffic to the specified PCP server.

```
[edit services service-set service-set-name ]
user @host# set pcp-rule rule-name | rule-listname
```



NOTE: Your service set must also identify any required **nat-rule** and **software-rule**.

SYSLOG Message Configuration

A new syslog class, configuration option, **pcp-logs**, has been provided to control PCP log generation. It provides the following levels of logging:

- **protocol**—All logs related to mapping creation, deletion are included at this level of logging.
- **protocol-error**—All protocol error related logs (such as mapping refresh failed, PCP look up failed, mapping creation failed). are included in this level of logging.
- **system-error**—Memory and infrastructure errors are included in this level of logging.

Example: Configuring Port Control Protocol with NAPT44

- [Requirements on page 170](#)
- [Overview on page 170](#)
- [PCP Configuration on page 170](#)

Requirements

Hardware Requirements

- UEs with PCP clients.
- An MX 3D Router with an MS-DPC services PIC.


Software Requirements

- Junos OS 13.2
- Layer-3 Services Package

Overview

An ISP wants to enable UEs with PCP clients to maintain connections to servers without timing out. The PCP clients generate PCP requests for the type and duration of the connection they require. Connections may be of a long duration, such as applications using a webcam, or a shorter duration, such as online games. An MX 3D router provides a PCP server to interpret PCP client requests, and NAPT44. [Figure 12 on page 170](#) shows the basic topology for this example.

Figure 12: PCP with NAPT44



PCP Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

```

set chassis fpc 2 pic 0 adaptive-services service-package layer-3
set interfaces sp-2/0/0 services-options inactivity-timeout 180 cgn-pic
set interfaces sp-2/0/0 unit 0 family inet
  
```

```

set interfaces xe-3/2/0 unit 0 family inet service input service-set sset_0
set interfaces xe-3/2/0 unit 0 family inet service output service-set sset_0
set interfaces xe-3/2/0 unit 0 family inet address 30.0.0.1/24
set interfaces xe-5/0/0 unit 0 family inet address 25.0.0.1/24
set services nat pool pcp-pool address 44.0.0.0/16
set services nat pool pcp-pool port automatic random-allocation address-allocation
    round-robin
set services nat pool pcp-pool address-allocation round-robin
set services nat rule pcp-rule match-direction input
set services nat rule pcp-rule term t0 then translated source-pool pcp-pool
    translation-type napt-44
set services nat rule pcp-rule term t0 then translated mapping-type endpoint-independent
    filtering-type endpoint-independent
set services nat rule pcp-rule term t0 then translated mapping-type endpoint-independent
    filtering-type endpoint-independent
set services pcp server pcp-s1 ipv4-address 124.124.124.122 mapping-lifetime-minimum
    600 mapping-lifetime-maximum 600
set services pcp server pcp-s1 mapping-lifetime-minimum 600
    mapping-lifetime-maximum 86500
set services pcp server pcp-s1 short-lifetime-error 120 long-lifetime-error 1200
set services pcp server pcp-s1 max-mappings-per-client 128 pcp-options third-party
    prefer-failure
set services service-set sset_0 pcp-rules r1
set services service-set sset_0 nat-rules pcp-rule
set services service-set sset_0 interface-service service-interface sp-2/0/0.0

```

Chassis Configuration

Step-by-Step Procedure

To configure the service PIC (FPC 2 Slot 0) with the Layer 3 service package:

1. Go to the [edit chassis] hierarchy level.

```
user@host# edit chassis
```
2. Configure the Layer 3 service package.

```
[edit chassis]
user@host# set fpc 2 pic 0 adaptive-services service-package layer-3
```

Results `user@host# show chassis fpc 2 pic 0`

```

pcp-rules pcp-napt44-rule;
nat-rules pcp-rule;
interface-service {
    service-interface sp-2/0/0.0;
}

```

Interface Configuration

Step-by-Step Procedure

1. Configure the services MS-DPC.

```
user@host# set interfaces sp-2/0/0 services-options inactivity-timeout 180 cgn-pic
user@host# set interfaces sp-2/0/0 unit 0 family inet
```
2. Configure the customer-facing interface used for NAT and PCP services.

```
user@host# set interfaces xe-3/2/0 unit 0 family inet service input service-set
sset_0
user@host# set interfaces xe-3/2/0 unit 0 family inet service output service-set
sset_0
user@host# set interfaces xe-3/2/0 unit 0 family inet address 30.0.0.1/24
```

3. Configure the Internet-facing interface.

```
user@host# set interfaces xe-5/0/0 unit 0 family inet address 25.0.0.1/24
```

Results

```
user@host#
sp-2/0/0 {
  services-options {
    inactivity-timeout 180;
    cgn-pic;
  }
  unit 0 {
    family inet;
  }
}
xe-3/2/0 {
  unit 0 {
    family inet {
      service {
        input {
          service-set sset_0;
        }
        output {
          service-set sset_0;
        }
      }
      address 30.0.0.1/24;
    }
  }
}
xe-5/0/0 {
  unit 0 {
    family inet {
      address 25.0.0.1/24;
    }
  }
}
```

NAT Configuration

Step-by-Step Procedure

1. Go the **[edit services nat]** hierarchy.

```
user@host# edit services nat
```
2. Configure a NAT pool called **pcp-pool**.

```
[edit services nat]
user@host# set pool pcp-pool address 44.0.0.0/16
user@host# set pool pcp-pool port automatic random-allocation
user@host# set pool pcp-pool address-allocation round-robin
```
3. Configure a NAT rule called **pcp-rule**.

```
[edit services nat]
```

```

user@host# set rule pcp-rule term t0 then translated source-pool pcp-pool
translation-type napt-44
user@host# set rule pcp-rule term t0 then translated mapping-type
endpoint-independent filtering-type endpoint-independent

```

Results

```

user@host# show services nat
pool pcp-pool {
  address 44.0.0.0/16;
  port {
    automatic {
      random-allocation;
    }
  }
  address-allocation round-robin;
}
rule pcp-rule {
  match-direction input;
  term t0 {
    then {
      translated {
        source-pool pcp-pool;
        translation-type {
          napt-44;
        }
        mapping-type endpoint-independent;
        filtering-type {
          endpoint-independent;
        }
      }
    }
  }
}

```

PCP Configuration

Step-by-Step Procedure To configure the PCP server and PCP rule options.

1. Go to the **edit services pcp** hierarchy level for server **pcp-s1**

```

user@host# edit services pcp server pcp-s1

```
2. Configure the PCP server options.


```

[edit services pcp server pcp-s1]
user@host# set ipv4-address 124.124.124.122
user@host# set mapping-lifetime-minimum 600
user@host# set mapping-lifetime-maximum 86500
user@host# set short-lifetime-error 120
user@host# set long-lifetime-error 1200
user@host# set max-mappings-per-client 128
user@host# set pcp-options third-party prefer-failure

```
3. Create the PCP rule.


```

[edit services pcp rule pcp-napt44-rule
user@host# edit rule pcp-napt44-rule

```
4. Configure the PCP rule options.

```
[edit services pcp rule pcp-napt44-rule]
user@host# set match-direction input
user@host# set term t0 then pcp-server pcp-s1
```

Results regress@montag# show services pcp

```
server pcp-s1 {
  ipv4-address 124.124.124.122;
  mapping-lifetime-minimum 600;
  mapping-lifetime-maximum 86500;
  short-lifetime-error 120;
  long-lifetime-error 1200;
  max-mappings-per-client 128;
  pcp-options third-party prefer-failure;
}
rule pcp-napt44-rule {
  match-direction input;
  term t0 {
    then {
      pcp-server pcp-s1;
    }
  }
}
```

Service Set Configuration

Step-by-Step Procedure

1. Create a service set, **sset_0**, at the **edit services service-set** hierarchy level.

```
user@router# edit services service-set sset_0

service-set sset_0 {
  pcp-rules pcp-napt44-rule;
  nat-rules pcp-rule;
  interface-service {
    service-interface sp-2/0/0.0;
  }
}
```

2. Identify the NAT rule associated with the service set.

```
[edit services service-set sset_0]
user@router# set nat-rules pcp-rule
```

3. Identify the PCP rule associated with the service set.

```
[edit services service-set sset_0]
user@router# set pcp-rules r1
```

4. Identify the service interface associated with the service set.

```
[edit services service-set sset_0]
user@router# set interface-service service-interface sp-2/0/0.0
```


Results user@host# show
pcp-rules pcp-napt44-rule;
nat-rules pcp-rule;
interface-service {
 service-interface sp-2/0/0.0;
}

Automatically Assigning Ports Using Port Block Allocation

- [Secured Port Block Allocation for NAPT on page 177](#)
- [Interim Logging for Port Block Allocation on page 177](#)
- [Configuring Secured Port Block Allocation on page 178](#)
- [Configuring Deterministic Port Block Allocation on page 180](#)

Secured Port Block Allocation for NAPT

Secured port block allocation can be used for translation types **napt-44** and **stateful-nat64**.

When allocating blocks of ports, the most recently allocated block is the current active block. New requests for NAT ports are served from the active block. Ports are allocated randomly from the current active block.

When you configure secured port block allocation, you can specify the following:

- **block-size**
- **max-blocks-per-address**
- **active-block-timeout**

Related Documentation

- [Configuring Secured Port Block Allocation on page 178](#)

Interim Logging for Port Block Allocation

With port block allocation we generate one syslog log per set of ports allocated for a subscriber. These logs are UDP based and can be lost in the network, particularly for long-running flows. Interim logging triggers re-sending the above logs at a configured interval for active blocks that have traffic on at least one of the ports of the block.

Interim logging is activated by including the **pba-interim-logging-interval** statement under **services-options**.

- Related Documentation**
- [Configuring NAT Session Logs on page 209](#)
 - [Secured Port Block Allocation for NAT on page 177](#)

Configuring Secured Port Block Allocation

To configure secured port block allocation:

1. At the **[edit services nat pool *poolname*]** hierarchy level, create a pool.

```
user@host# edit services nat pool poolname
```

For example:

```
user@host# edit services nat pool pba-pool1
```

2. Define the range of addresses to be translated, specifying the upper and lower limits of the range or an address prefix that describes the range.

```
[edit services nat pool pba-pool1]  
user@host# set address-range low address high address
```

Or

```
user@host# set address address-prefix
```

For example:

```
[edit services nat pool pba-pool1]  
user@host# set address 203.0.113.0/24
```

3. Define the range of ports to be used in the translation, or use automatic port assignment by the Junos OS. You can optionally specify random assignment of ports (sequential assignment is the default).

```
[edit services nat pool pba-pool1]  
user@host# set port range low address high address random
```

Or

```
user@host# set port automatic random-allocation
```

For example:

```
[edit services nat pool pba-pool1]  
user@host# set port range low 256 high 511 random
```

Or

```
[edit services nat pool pba-pool1]  
user@host# set port automatic random-allocation
```



NOTE: When you configure a port range, the range should be a multiple of the port block-size value (see Step 4). When the NAT pool port range is *not* a multiple of the port block-size value, the number of ports or port-blocks that are effectively available for use is less than the configured number of ports and port-blocks. The port block allocation mechanism uses ports in the range 0 through 1023 of a NAT address.

When you configure automatic assignment of ports, the available port range for allocation is 1024 through 65535. Automatic allocation can result in no ports being available for use. Use the `show services nat pool` command on the Routing Engine after you configure the port block allocation method to determine the number of ports and port blocks available for allocation to users.

4. Configure secured port block allocation. Specify **active-block-timeout**, **block-size**, and **max-blocks-per-address**, or accept the default values for those options.

```
[edit services nat pool pba-pool1]
user@host# set secured-port-block-allocation active-block-timeout
active-block-timeout block-size block-size max-blocks-per-address
max-blocks-per-address
```

For example:

```
[edit services nat pool pba-pool1]
user@host# set secured-port-block-allocation active-block-timeout 120 block-size
256 max-blocks-per-address 12
```



NOTE: In order for secured-port-block-allocation configuration changes to take effect, you must reboot the services PIC whenever you change any of the following nat pool options:

- *pool-name*
- address or address-range
- port range
- port secured-port-block-allocation block-size
- port secured-port-block-allocation max-blocks-per-address.
- port secured-port-block-allocation active-block-timeout.
- from hierarchy in the nat rule



NOTE: MS-MICs and MS-MPCs support up to a maximum of nine million port blocks per NPU. If your configuration exceeds this maximum supported number, one or more service sets might not be activated on that NPU.

Related Documentation • [Network Address Translation Configuration Overview on page 65](#)

Configuring Deterministic Port Block Allocation

To configure deterministic port block allocation:

1. At to the `[edit services nat pool poolname]` hierarchy level, create a pool.

```
user@host# edit services nat pool poolname
```

For example:

```
user@host# edit services nat pool pba-pool2
```

2. Define the range of addresses to be translated, specifying the upper and lower limits of the range or an address prefix that describes the range.

```
[edit services nat pool pba-pool1]
user@host# set address-range low address high address
```

Or

```
user@host# set address address-prefix
```

For example:

```
[edit services nat pool pba-pool1]
user@host# set address-range low 32.32.32.1 high 32.32.32.253
```

3. Specify automatic port assignment by the Junos OS.

```
[edit services nat pool pba-pool1]
user@host# set port automatic sequential
```



NOTE: Starting with Junos OS release 14.2, the `sequential` option is introduced to enable you to configure sequential allocation of ports. The `sequential` and `random-allocation` options available with the `port automatic` statement at the `[edit services nat pool nat-pool-name]` hierarchy level are mutually exclusive. You can include the `sequential` option for sequential allocation and the `random-allocation` option for random delegation of ports. By default, sequential allocation of ports takes place if you include only the `port automatic` statement at the `[edit services nat pool nat-pool-name]` hierarchy level. The `auto` option is hidden and is deprecated in Junos OS Release 14.2 and later, and is only maintained for backward compatibility. It might be removed completely in a future software release.

4. Configure deterministic port block allocation. Specify `block-size` or accept the default value of 512.

. You can also specify `include-boundary-addresses` if you want the lowest and highest addresses in the source address range of a NAT rule to be translated when the NAT pool is used.

```
[edit services nat pool pba-pool1]
```

```
user@host# set port deterministic-port-block-allocation block-size block-size  
include-boundary-addresses
```

For example:

```
[edit services nat pool pba-pool1]  
user@host# set port deterministic-port-block-allocation block-size 256
```



NOTE: In order for deterministic-port-block-allocation configuration changes to take effect, you must reboot the services PIC whenever you change any of the following nat pool options:

- address or address-range
- port range
- port deterministic-port-block-allocation block-size

**Related
Documentation**

- [Network Address Translation Configuration Overview on page 65](#)

Connecting Specific Ports and Addresses Using Port Forwarding

- [Configuring Port Forwarding for Static Destination Address Translation on page 183](#)
- [Configuring Port Forwarding Without Destination Address Translation on page 186](#)
- [Example: Configuring Port Forwarding with Twice NAT on page 187](#)

Configuring Port Forwarding for Static Destination Address Translation

Starting with Junos OS Release 11.4, you can map an external IP address and port with an IP address and port in a private network. This allows the destination address and port of a packet to be changed to reach the correct host in a Network Address Translation (NAT) gateway. The translation facilitates reaching a host within a masqueraded, typically private, network, based on the port number on which the packet was received from the originating host. Port forwarding allows remote computers, such as public machines on the Internet, to connect to a non-standard port (port other than 80) of a specific computer within a private network. An example of this type of destination is the host of a public HTTP server within a private network. Port forwarding is supported only with **dnat-44** and **twice-napt-44** on IPv4 networks. Port forwarding works only with the FTP application-level gateway (ALG). Port forwarding also supports endpoint-independent mapping (EIM), endpoint-independent filtering (EIF), and address pooling paired (APP). Port forwarding has no support for technologies such as IPv6 rapid deployment (6rd) and dual-stack lite (DS-Lite) that offer IPv6 services over IPv4 infrastructure.

To configure destination address translation with port forwarding in IPv4 networks:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

```
[edit]
user@host# edit services nat
```

2. Configure the NAT pool with an address.

```
[edit services nat]
user@host# set pool pool-name address address
```

In the following example, **dest-pool** is used as the pool name and **4.1.1.2** as the address.

```
user@host# set pool dest-pool address 4.1.1.2
```

3. Configure the rule, match direction, term, and destination address.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction term term-name from
destination-address address
```

In the following example, the name of the rule is **rule-dnat44**, the match direction is **input**, the name of the term is **t1**, and the address is **20.20.20.20**.

```
[edit services nat]
user@host# set rule rule-dnat44 match-direction input term t1 from
destination-address 20.20.20.20
```

4. Configure the destination port range.

```
[edit services nat]
user@host# set rule rule-dnat44 match-direction match-direction term term-name from
destination-port range range high | low
```

In the following example, the upper port range is **50** and the lower port range is **20**.

```
[edit services nat]
user@host# set rule rule-dnat44 match-direction input term t1 from destination-port
range range high 50 low 20
```

5. Go to the **[edit services nat rule rule-dnat44 term t1]** hierarchy level.

```
[edit services nat]
user@host# edit rule rule-dnat44 term t1
```

6. Configure the destination pool.

```
[edit services nat rule rule-dnat44 term t1]
user@host# set then translated destination-pool dest-pool-name
```

In the following example, the destination pool name is **dest-pool**, and the translation type is **dnat-44**.

```
[edit services nat rule rule-dnat44 term t1]
user@host# set then translated destination-pool dest-pool
```

7. Configure the mapping for port forwarding and the translation type.

```
[edit services nat rule rule-dnat44 term t1]
user@host# set then port-forwarding-mappings map-name translation-type
translation-type
```

In the following example, the port forwarding map name is **map1**, and the translation type is **dnat-44**.

```
[edit services nat rule rule-dnat44 term t1]
user@host# set then port-forwarding-mappings map1 translation-type dnat-44
```

8. Go to the **[edit services nat port-forwarding map1]** hierarchy level.

```
[edit services nat]
user@host# edit port-forwarding map1
```

9. Configure the mapping for port forwarding.

```
[edit port-forwarding map1]
user@host# set destined-port port-id
user@host# set translated-port port-id
```

In the following example, the destination port is **45** and the translated port is **23**.

```
[edit port-forwarding map1]
user@host# set destined-port 23
user@host# set translated-port 45
```

**NOTE:**

- Multiple port mappings are supported with port forwarding. Up to 32 port maps can be configured for port forwarding.
- The destination port should not overlap the port range configured for NAT.

10. Verify the configuration by using the **show** command at the **[edit services nat]** hierarchy level.

```
[edit services]
user@host# show
nat {
  pool dest-pool {
    address 4.1.1.2/32;
  }
  rule rule-dnat44 {
    match-direction input;
    term t1 {
      from {
        destination-address {
          20.20.20.20/32;
        }
        destination-port {
          range low 20 high 50;
        }
      }
      then {
        port-forwarding-mappings map1;
        translated {
          destination-pool dest-pool;
          translation-type {
            dnat-44;
          }
        }
      }
    }
  }
}
port-forwarding map1 {
  destined-port 45;
  translated-port 23;
}
```

**NOTE:**

- A similar configuration is possible with twice NAT for IPv4. See [“Example: Configuring Port Forwarding with Twice NAT”](#) on page 187.
- Port forwarding and stateful firewall can be configured together. Stateful firewall has precedence over port forwarding.

Related Documentation • [Configuring Static Destination Address Translation in IPv4 Networks on page 111](#)

Configuring Port Forwarding Without Destination Address Translation

Starting with Junos OS Release 12.1, you can configure port forwarding without translating a destination address.

To configure port forwarding without destination address translation in IPv4 networks:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

```
[edit]
user@host# edit services nat
```

2. Configure the rule, match direction, term, and destination address.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction term term-name
```

In the following example, the name of the rule is **rule-port-forwarding**, the match direction is **input**, and the name of the term is **t1**.

```
[edit services nat]
user@host# set rule rule-port-forwarding match-direction input term t1
```

3. Go to the **[edit services nat rule rule-port-forwarding term t1]** hierarchy level.

```
[edit services nat]
user@host# edit rule rule-port-forwarding term t1
```

4. Specify that there is no address translation for this rule.

```
[edit services nat rule rule-port-forwarding term t1]
user@host# set then no-translation
```

5. Configure the mapping for port forwarding and the translation type.

```
[edit services nat rule rule-port-forwarding term t1]
user@host# set then port-forwarding-mappings map-name
```

In the following example, the port forwarding map name is **map1**.

```
[edit services nat rule rule-port-forwarding term t1]
user@host# set then port-forwarding-mappings map1
```

6. Go to the **[edit services nat port-forwarding map1]** hierarchy level.

```
[edit services nat]
user@host# edit port-forwarding map1
```

7. Configure the mapping for port forwarding.

```
[edit port-forwarding map1]
user@host# set destined-port port-id
user@host# set translated-port port-id
```

In the following example, the destination port is **45** and the translated port is **23**.

```
[edit port-forwarding map1]
user@host# set destined-port 23
user@host# set translated-port 45
```

**NOTE:**

- Multiple port mappings are supported with port forwarding. Up to 32 port maps can be configured for port forwarding.
- The destination port should not overlap the port range configured for NAT.

8. Verify the configuration by using the **show** command at the **[edit services nat]** hierarchy level.

```
[edit services]
user@host# show
nat {
  rule rule-port-forwarding {
    match-direction input;
    term t1 {
      then {
        port-forwarding-mappings map1;
        no-translation      }
      }
    }
  }
  port-forwarding map1 {
    destined-port 45;
    translated-port 23;
  }
}
```



NOTE: Port forwarding and stateful firewall can be configured together. Stateful firewall has precedence over port forwarding.

Example: Configuring Port Forwarding with Twice NAT

The following example configures port forwarding with **twice-napt-44** as the translation type. The example also has stateful firewall and multiple port maps configured.

```
[edit services]
user@host# show
service-set in {
  syslog {
    host local {
      services any;
    }
  }
  stateful-firewall-rules r;
  nat-rules r;
  interface-service {
    service-interface sp-10/0/0.0;
  }
}
stateful-firewall {
  rule r {
    match-direction input;
    term t {
```

```
        from {
            destination-port {
                range low 20 high 5000;
            }
        }
        then {
            reject;
        }
    }
}
nat {
    pool x {
        address 12.0.0.2/32;
    }
    rule r {
        match-direction input;
        term t {
            from {
                destination-address {
                    14.0.0.2/32;
                }
                destination-port {
                    range low 10 high 20000;
                }
            }
            then {
                port-forwarding-mappings y;
                translated {
                    destination-pool x;
                    translation-type {
                        twice-napt-44;
                    }
                }
            }
        }
    }
}
port-forwarding y {
    destined-port 45;
    translated-port 23;
    destined-port 55;
    translated-port 33;
    destined-port 65;
    translated-port 43;
}
}
adaptive-services-pics {
    traceoptions {
        file sp-trace;
        flag all;
    }
}
```



NOTE:

- Stateful firewall has precedence over port forwarding. In this example, for instance, no traffic destined to any port between 20 and 5000 will be translated.
- Up to 32 port maps can be configured.

**Related
Documentation**

- [Configuring Port Forwarding for Static Destination Address Translation on page 183](#)

Allocating a Few Public Addresses to Many Private Hosts Using Dynamic NAT

- [Configuring Dynamic Address-Only Source Translation in IPv4 Networks on page 191](#)
- [Example: Dynamic Source NAT as a Next-Hop Service on page 195](#)
- [Example: Assigning Addresses from a Dynamic Pool for Static Use on page 197](#)

Configuring Dynamic Address-Only Source Translation in IPv4 Networks

In IPv4 networks, dynamic address translation (dynamic NAT) is a mechanism to dynamically translate the destination traffic without port mapping. To use dynamic NAT, you must specify a source pool name, which includes an address configuration.

To configure dynamic NAT in IPv4 networks:

1. In configuration mode, go to the **[edit services]** hierarchy level.

```
[edit]
user@host# edit services
```

2. Configure the service set and NAT rule.

```
[edit services]
user@host# set service-set service-set-name nat-rules rule-name
```

In the following example, the name of the service set is **s1**, and the name of the NAT rule is **rule-dynamic-nat44**.

```
[edit services]
user@host# set service-set s1 nat-rules rule-dynamic-nat44
```

3. Go to the **[interface-service]** hierarchy level for the service set.

```
[edit services]
user@host# edit service-set s1 interface-service
```

4. Configure the service interface.

```
[edit services service-set s1 interface-service]
user@host# set service-interface service-interface-name
```

In the following example, the name of the service interface is **ms-0/1/0**.



NOTE: If the service interface is not present in the router, or the specified interface is not functional, the following command can result in an error.

```
[edit services service-set s1 interface-service]
user@host# set service-interface ms-0/1/0
```

5. Go to the **[edit services nat]** hierarchy level. Issue the following command from the top of the services hierarchy, or use the **top** keyword.

```
[edit services service-set s1 interface-service]
user@host# top edit services nat
```

6. Configure the NAT pool with an address.

```
[edit services nat]
user@host# set pool pool-name address address
```

In the following example, the name of the pool is **source-dynamic-pool**, and the address is **10.10.10.0**.

```
[edit services nat]
user@host# set pool source-dynamic-pool address 10.10.10.0
```

7. Configure the rule, match direction, term, and source address.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction term term-name from
source-address address
```

In the following example, the name of the rule is **rule-dynamic-nat44**, the match direction is **input**, the name of the term is **t1**, and the source address is **3.1.1.0**.

```
[edit services nat]
user@host# set rule rule-dynamic-nat44 match-direction input term t1 from
source-address 3.1.1.0
```

8. Go to the **[edit rule rule-dynamic-nat-44 term t1]** hierarchy level.

```
[edit services nat]
user@host# edit rule rule-dynamic-nat44 term t1
```

9. Configure the source pool and the translation type.

```
[edit services nat rule rule-dynamic-nat44 term t1]
user@host# set then translated source-pool src-pool-name translation-type
translation-type
```

In the following example, the name of the source pool is **source-dynamic-pool** and the translation type is **dynamic-nat44**.

```
[edit services nat rule rule-dynamic-nat44 term t1]
user@host# set then translated source-pool source-dynamic-pool translation-type
dynamic-nat44
```

10. Go to the **[edit services adaptive-services-pics]** hierarchy level. In the following command, the **top** keyword ensures that the command is run from the top of the hierarchy.

```
[edit services nat rule rule-dynamic-nat44 term t1]
```

```
user@host# top edit services adaptive-services-pics
```

11. Configure the trace options.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag tracing parameter
```

In the following example, the tracing parameter is configured as **all**.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag all
```

12. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
service-set s1 {
  nat-rules rule-dynamic-nat44;
  interface-service {
    service-interface ms-0/1/0;
  }
}
nat {
  pool source-dynamic-pool {
    address 10.1.1.0/24;
  }
  rule rule-dynamic-nat44 {
    match-direction input;
    term t1 {
      from {
        source-address {
          3.1.1.0/24;
        }
      }
      then {
        translated {
          destination-pool source-dynamic-pool;
          translation-type {
            dynamic-nat44;
          }
        }
      }
    }
  }
}
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```

The following example configures the translation type as **dynamic-nat44**.

```
[edit services]
user@host# show
service-set s1 {
  nat-rules rule-dynamic-nat44;
  interface-service {
    service-interface ms-0/1/0;
  }
}
```

```
nat {
  pool source-dynamic-pool {
    address 10.1.1.0/24;
  }
  rule rule-dynamic-nat44 {
    match-direction input;
    term t1 {
      from {
        source-address {
          3.1.1.0/24;
        }
      }
      then {
        translated {
          destination-pool source-dynamic-pool;
          translation-type {
            dynamic-nat44;
          }
        }
      }
    }
  }
}
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```

The following configuration specifies that NAT is not performed on incoming traffic from the source address **192.168.20.24/32** by providing a NAT rule term **t0** that configures **no-translation**. Dynamic NAT is performed on all other incoming traffic, as configured by term **t1** of the NAT rule.

```
[edit services nat]
pool my-pool {
  address-range low 10.10.10.1 high 10.10.10.16;
  port automatic;
}
rule src-nat {
  match-direction input;
  term t0 {
    from {
      source-address 192.168.20.24/32;
    }
    then {
      no-translation;
    }
  }
  term t1 {
    then {
      translated {
        translation-type dynamic-nat44;
        source-pool my-pool;
      }
    }
  }
}
```

The following configuration performs NAT using the source prefix **20.20.10.0/24** without defining a pool.

```
[edit services nat]
rule src-nat {
  match-direction input;
  term t1 {
    then {
      translation-type dynamic-nat44;
      source-prefix 20.20.10.0/24;
    }
  }
}
```

The following configuration performs NAT using the destination prefix **20.20.10.0/32** without defining a pool.

```
[edit services nat]
rule src-nat {
  match-direction input;
  term t1 {
    from {
      destination-address 10.10.10.10/32;
    }
    then {
      translation-type dnat44;
      destination-prefix 20.20.10.0/24;
    }
  }
}
```

Example: Dynamic Source NAT as a Next-Hop Service

The following example shows dynamic-source NAT applied as a next-hop service:

```
[edit interfaces]
ge-0/2/0 {
  unit 0 {
    family mpls;
  }
}
sp-1/3/0 {
  unit 0 {
    family inet;
  }
  unit 20 {
    family inet;
  }
  unit 32 {
    family inet;
  }
}
[edit routing-instances]
protected-domain {
  interface ge-0/2/0.0;
  interface sp-1/3/0.20;
```

```
instance-type vrf;
route-distinguisher 10.58.255.17:37;
vrf-import protected-domain-policy;
vrf-export protected-domain-policy;
routing-options {
    static {
        route 0.0.0.0/0 next-hop sp-1/3/0.20;
    }
}
[edit policy-options]
policy-statement protected-domain-policy {
    term t1 {
        then reject;
    }
}
[edit services]
stateful-firewall {
    rule allow-all {
        match-direction input;
        term t1 {
            then {
                accept;
            }
        }
    }
}
nat {
    pool my-pool {
        address 10.58.16.100;
        port automatic;
    }
    rule hide-all {
        match-direction input;
        term t1 {
            then {
                translated {
                    source-pool my-pool;
                    translation-type napt-44;
                }
            }
        }
    }
}
service-set null-sfw-with-nat {
    stateful-firewall-rules allow-all;
    nat-rules hide-all;
    next-hop-service {
        inside-service-interface sp-1/3/0.20;
        outside-service-interface sp-1/3/0.32;
    }
}
```

Example: Assigning Addresses from a Dynamic Pool for Static Use

The following configuration statically assigns a subset of addresses that are configured as part of a dynamic pool (**dynamic-pool**) to two separate static pools (**static-pool** and **static-pool2**).

```
[edit services nat]
pool dynamic-pool {
  address 20.20.10.0/24;
}
pool static-pool {
  address-range low 20.20.10.10 high 10.20.10.12;
}
pool static-pool2 {
  address 20.20.10.15/32;
}
rule src-nat {
  match-direction input;
  term t1 {
    from {
      source-address 30.30.30.0/24;
    }
    then {
      translation-type dynamic-nat44;
      source-pool dynamic-pool;
    }
  }
  term t2 {
    from {
      source-address 10.10.10.2;
    }
    then {
      translation-type basic-nat44;
      source-pool static-pool;
    }
  }
  term t3 {
    from {
      source-address 10.10.10.10;
    }
    then {
      translation-type basic-nat44;
      source-pool static-pool2;
    }
  }
}
```


Achieving Line-Rate, Low-Latency Translations Using Inline NAT

- [Inline Network Address Translation Overview for MPC Types 1, 2, and 3 on page 199](#)
- [Example: Configuring Inline Network Address Translation - Interface-Service Service Set on page 201](#)

Inline Network Address Translation Overview for MPC Types 1, 2, and 3

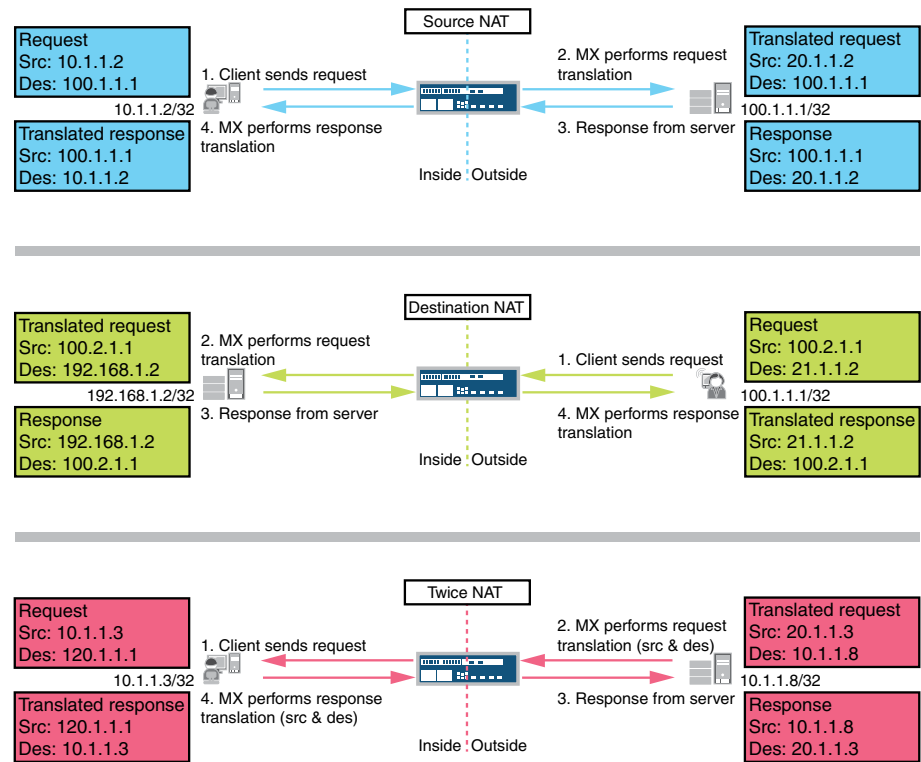
Inline network address translation (NAT) uses the capabilities of the Modular Port Concentrator (MPC) line card, eliminating the need for a MultiServices Dense Port Concentrator (MS-DPC) for NAT. Consequently, you can achieve line-rate, low-latency address translations (up to 120 Gbps per slot). The current implementation provides:

- 1:1 static address mapping
- Bidirectional mapping - source NAT for outbound traffic and destination NAT for inbound traffic
- No limit on number of flows
- Support for Source, destination, and twice NAT, as shown in [Figure 13 on page 200](#)



NOTE: Inline NAT is generally only the `basic-nat44` translation type, and implements destination NAT and twice NAT by applying NAT at the egress interface or to back-to-back, as shown in the following figure.

Figure 13: Supported Inline NAT Types



g041381

To configure inline NAT, you define your service interface as type **si-** (service-inline) interface. You must also reserve adequate bandwidth for the inline interface. This enables you to configure both interface or next-hop service-sets used for NAT. The **si-** interface serves as a “virtual service PIC”.



NOTE: Only static source NAT is supported. Port translation and dynamic NAT are not supported. An MS-DPC or MS-PIC will still be needed for any stateful-firewall processing.

Related Documentation

- [Network Address Translation Configuration Overview on page 65](#)
- [Example: Configuring Inline Network Address Translation - Interface-Service Service Set on page 201](#)
- [Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 61](#)

Example: Configuring Inline Network Address Translation - Interface-Service Service Set

- [Requirements on page 201](#)
- [Overview on page 201](#)
- [Configuration on page 203](#)

Requirements

This example uses the following hardware and software components:

- MX-series router
- Modular Port Concentrator (MPC) with Trio chipset
- Junos OS Release 11.4R1 or higher

Overview

This example is configured for the network of a large financial services firm. This Application Service Provider (ASP) has an IP/MPLS-based backbone and provides L3VPN connectivity. In our example, the ASP acts like an Internet Service Provider (ISP) and its servers have public IPv4 addresses.

A large subscriber base relies heavily on the market data feeds that the ASP provides. Like many of the enterprise networks today, a private addressing scheme has been in place for majority of ASP's customers. NAT is required to maintain access to ASP's shared services.

Requirements for the solution include:

- Ease subscriber addressing challenges of their by providing NAT services in ASP's network.
- Support access to common services by a large number of customers, even when these are hosted across in different VRFs and use overlapping addresses.
- Provide high throughput, low latency packet forwarding with NAT enabled.
- Provide operational simplicity and efficiency.
- Reduce cost of operations.

By deploying Juniper's MX's inline NAT service, the ASP can offer scalable solutions with uncompromised performance that fit the requirements of financial markets customers. Operational cost can be dramatically reduced by eliminating the need for a dedicated services PIC. Enabling subscribers to keep their existing addressing scheme by outsourcing the address translation function to the ASP greatly simplifies their network operations.


Topology

The topology for this application is show in [Figure 14 on page 202](#)

The ASP's shared services are located on LAN segment 101.1.1.0/24 behind PE3. PE1 and PE2 are used to connect subscribers. Traditional MPLS-VPN is deployed between provider edge routers. In the case of PE1, subscriber A and B have overlapping addressing schemes of 10.1.1.0/24; NAT is needed so the subscribers can access the same server. NAT pools 192.168.1.0/24 and 192.168.2.0/24 have been allocated to customer A and B respectively.

We will use host 10.1.1.2 from customer A to illustrate packet flow at a high level, as shown in Figure 14 on page 202

Figure 14: Deploy Inline NAT within L3VPN



1. CE1 forwards request from host 10.1.1.2 with a server destination of 101.1.1.2
2. With configured service set on PE1 for VRF_A, source address of 10.1.1.2 will be translated into 192.168.1.2. VPN label and IGP label will be imposed after the translation.
3. Packets will then be forwarded to PE3 using IGP label
4. PE3 receives the packet and performs a lookup in its VPN routing table. It then forwards the packets to server 101.1.1.2 after label disposition.
5. The server returns the packet with destination address of 192.168.1.2.
6. PE3 imposes VPN and IGP labels for the above destination and label switched the packets to PE1.
7. PE1 sends the packet to VRF_A after a FIB lookup. Destination address 192.168.1.2 will be translated 10.1.1.2.
8. CE1 receives the packets for host 10.1.1.2 and forwards them on.

Configuration

By using a **si-** (service-inline) interface, the operator can configure both **interface-service** and **next-hop** service-sets to perform inline NAT. This example uses the **interface-service** service set.

To configure inline NAT, perform these tasks:

- [Configure Interfaces on page 203](#)
- [Configuring Bandwidth for the Service Inline \(si-\) Interface on page 205](#)
- [Configuring NAT Pool and Rule on page 206](#)
- [Configuring the Service Set on page 207](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

```
set interfaces si-0/0/0 unit 0 family inet
set interfaces ge-1/0/0 unit 0 family inet service input service-set nat1
set interfaces ge-1/0/0 unit 0 family inet service output service-set nat1
set interfaces ge-1/0/0 unit 0 family inet address 100.1.1.252/30
set interfaces ge-1/0/0 unit 1 family inet service input service-set nat2
set interfaces ge-1/0/0 unit 1 family inet service output service-set nat2
set interfaces ge-1/0/0 unit 1 family inet address 100.1.1.250/30
set interfaces ge-3/0/0 unit 0 family inet service input service-set nat3
set interfaces ge-3/0/0 unit 0 family inet service output service-set nat3
set interfaces ge-3/0/0 unit 0 family inet address 100.1.1.252/30
set chassis fpc 0 pic 0 inline-services bandwidth 10g
set services nat pool p1 address 192.1.68.1.0/24
set services nat pool p2 address 192.1.68.2.0/24
set services nat rule r1 match-direction input
set services nat rule r1 term t1 from source-address 10.1.1.0/24
set services nat rule r1 term t1 then translated source-pool p1 translation-type basic-nat44
set services nat rule r2 match-direction input
set services nat rule r2 term t1 from source-address 10.1.3.0/24
set services nat rule r2 term t1 then translated source-pool p2 translation-type basic-nat44
set services service-set nat1 nat-rules r1
set services service-set nat1 interface-service service-interface si-0/0/0.0
set services service-set nat2 nat-rules r2
set services service-set nat2 interface-service service-interface si-0/0/0.0
```

Configure Interfaces

Step-by-Step Procedure

To configure interfaces required for inline NAT:

1. Configure the inline interface for NAT services.


```
user@host# edit interfaces si-0/0/0
[edit interfaces si-0/0/0]
user@host# set unit 0 family inet
```
2. Configure the interface for traffic to be handled by service set nat1

```
user@host# edit interfaces ge-1/0/0
[edit interfaces ge-1/0/0]
user@host# edit unit 0 family inet service
[edit unit 0 family inet service]
user@host# set input service-set nat1 output service-set nat1
user@host# set address 100.1.1.252/30
```

3. Configure the interface for traffic to be handled by service set nat2

```
user@host# edit interfaces ge-1/0/0
[edit interfaces ge-1/0/0]
user@host# edit unit 1 family inet service input service
[edit unit 0 family inet service]
user@host# set input service-set nat2 output service-set nat2
user@host# set address 100.1.1.250/30
```

4. Configure the interface for traffic to be handled by service set nat3

```
user@host# edit interfaces ge-3/0/0
[edit interfaces ge-3/0/0]
user@host# edit unit 0 family inet service input service
[edit unit 0 family inet service]
user@host# set input service-set nat3 output service-set nat3
user@host# set address 100.1.1.252/30
```

```

Results si-0/0/0 {
        unit 0 {
            family inet;
        }
    }
    ge-1/0/0 {
        unit 0 {
            family inet {
                service {
                    input {
                        service-set nat1;
                    }
                    output {
                        service-set nat1;
                    }
                }
            }
            address 100.1.1.252/30;
        }
    }
    ge-1/0/0 {
        unit 1 {
            family inet {
                service {
                    input {
                        service-set nat2;
                    }
                    output {
                        service-set nat2;
                    }
                }
            }
            address 100.1.1.250/30;
        }
    }
    ge-3/0/0 {
        unit 0 {
            family inet {
                service {
                    input {
                        service-set nat3;
                    }
                    output {
                        service-set nat3;
                    }
                }
            }
            address 100.1.1.252/30;
        }
    }
}

```

Configuring Bandwidth for the Service Inline (si-) Interface

Step-by-Step Procedure

1. Go to the configuration hierarchy for the fpc and pic used for inline NAT services.

```

user@host# edit chassis fpc 0 pic 0
[edit chassis fpc - pic 0]

```
2. Set the bandwidth for inline services.

```

[edit chassis fpc 0 pic 0]

```

```
user@host# set inline-services bandwidth 10g
```

Configuring NAT Pool and Rule

Step-by-Step Procedure

1. Go to the services NAT hierarchy.

```
user@host# edit services nat
```
2. Configure two NAT pools.

```
[edit services nat]  
user@host# set nat pool p1 address 192.168.1.0/24  
user@host# set nat pool p2 address 192.168.2.0/24
```
3. Configure NAT rule for source pool p1.

```
[edit services nat]  
user@host# set nat rule r1 match-direction input  
user@host# set nat rule r1 term t1 from source-address 10.1.1.0/24 then  
[nat pool r1 term t1 from source-address 10.1.1.0/24 then]  
user@host# set translated source-pool p1 translation-type basic-nat44
```
4. Configure NAT rule for source pool p2.

```
[edit services nat]  
user@host# set nat rule r2 match-direction input  
user@host# edit nat rule r2 term t1 from source-address 10.1.3.0/24 then  
[nat pool r1 term t1 from source-address 10.1.3.0/24 then]  
user@host# set translated source-pool p3 translation-type basic-nat44
```



```

Results user@host# edit services nat
user@host# show

pool p1 {
    address 192.168.1.0/24;
}
pool p2 {
    address 192.168.2.0/24;
}

rule r1 {
    match-direction input;
    term t1 {
        from {
            source-address {
                10.1.1.0/24;
            }
        }
        then {
            translated {
                source-pool p1;
                translation-type {
                    basic-nat44;
                }
            }
        }
    }
}

rule r2 {
    match-direction input;
    term t1 {
        from {
            source-address {
                10.1.3.0/24;
            }
        }
        then {
            translated {
                source-pool p2;
                translation-type {
                    basic-nat44;
                }
            }
        }
    }
}

```

Configuring the Service Set

Step-by-Step Procedure

1. Configure a service set using NAT rule r1, associated with NAT pool p1.


```

user@host# edit services service-set nat1
[edit services service-set nat1]
user@host# set nat rules r1
user@host# set interface-service service-interface si-0/0/0.0

```
2. Configure a service set using NAT rule r2, associated with NAT pool p2.


```

user@host# edit services service-set nat2
[edit services service-set nat1]

```

```
user@host# set nat rules r2
user@host# set interface-service service-interface si-0/0/0.0
```

Results

```
user@host# edit services service-set nat1
user@host# show
nat-rules r1;
interface-service {
    service-interface si-0/0/0.0;
}
```

```
user@host# edit services service-set nat2
user@host# show
nat-rules r2;
interface-service {
    service-interface si-0/0/0.0;
}
```

Related Documentation

- [Inline Network Address Translation Overview for MPC Types 1, 2, and 3 on page 199](#)

Monitoring NAT

- [Configuring NAT Session Logs on page 209](#)
- [Monitoring NAT Pool Usage on page 210](#)

Configuring NAT Session Logs

You can configure session logs for NAT from the CLI. By default, session open and close logs are produced. However, you can request that only one type of log be produced.

To configure NAT session logs:

1. Go to the `[edit services service-set service-set-name syslog host class classname]` hierarchy level.

```
user@host# edit services service-set service-set-name syslog host class classname
```
2. Configure NAT logging using the `nat-logs` configuration statement.

```
[edit services service-set service-set-name syslog host class classname]
user@host# set nat-logs.
```
3. Configure session logging using the `session-logs` statement. Open and close logs are produced by default. Specify `open` or `close` to produce only one type of log.

```
[edit services service-set service-set-name syslog host class classname]
user@host# set session-logs.
```

Or

```
[edit services service-set service-set-name syslog host class classname]
user@host# set session-logs open.
```

Or

```
[edit services service-set service-set-name syslog host class classname]
user@host# set session-logs close.
```
4. For NAT sessions that use secured port block allocation (PBA), enter the `pba-interim-logging interval` option.

```
[edit services service-set service-set-name syslog host class classname]
user@host# top.
[edit]
user@host# set interfaces interface-name service-options
pba-interim-logging-intervale.
```

- Related Documentation**
- [Configuring System Logging for Service Sets on page 47](#)
 - [Interim Logging for Port Block Allocation on page 177](#)

Monitoring NAT Pool Usage

Purpose Use the **show services nat pool detail** command to find global NAT statistics related to pool usage. This command is frequently used in conjunction with the **show services stateful-firewall statistics** command.

Action user@host# **show services nat pool detail**

```
Interface: ms-1/0/0, Service set: s1
  NAT pool: dest-pool, Translation type: DNAT-44
    Address range: 10.10.10.2-10.10.10.2
  NAT pool: napt-pool, Translation type: NAPT-44
    Address range: 50.50.50.1-50.50.50.254
    Port range: 1024-63487, Ports in use: 0, Out of port errors: 0, Max ports
used: 0
  NAT pool: source-dynamic-pool, Translation type: DYNAMIC NAT44
    Address range: 40.40.40.1-40.40.40.254
    Out of address errors: 0, Addresses in use: 0
  NAT pool: source-static-pool, Translation type: BASIC NAT44
    Address range: 30.30.30.1-30.30.30.254
```

- Related Documentation**
- *NAT Objects MIB in SNMP MIBs and Traps Reference*
 - *Network Address Translation Resources—Monitoring MIB in SNMP MIBs and Traps Reference*
 - *Juniper Networks Enterprise-Specific NAT Traps on SRX Series Services Gateways in SNMP MIBs and Traps Reference*

PART 4

Transitioning to IPv6 Using Softwires

- [Softwires Overview on page 213](#)
- [Softwires Configuration Overview on page 219](#)
- [Transitioning to IPv6 Using 6to4 Softwires on page 223](#)
- [Transitioning to IPv6 Using DS-Lite Softwires on page 227](#)
- [Transitioning to IPv6 Using 6rd Softwires on page 245](#)
- [Monitoring and Troubleshooting Softwires on page 271](#)

CHAPTER 18

Softwires Overview

- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 213](#)

Tunneling Services for IPv4-to-IPv6 Transition Overview

The Junos OS enables service providers to transition to IPv6 by using softwire encapsulation and decapsulation techniques. A softwire is a tunnel that is created between softwire Customer Premises Equipment (CPE). A softwire CPE can share a unique common internal state for multiple softwires, making it a very light and scalable solution. When you use softwires, you need not maintain an interface infrastructure for each softwire, unlike a typical mesh of generic routing encapsulation (GRE) tunnels that would require you to do so. A softwire initiator at the customer end encapsulates native packets and tunnels them to a softwire concentrator at the service provider. The softwire concentrator decapsulates the packets and sends them to their destination. A softwire is created when a softwire concentrator receives the first tunneled packet of a flow and prepares for flow processing. The softwire exists as long as the softwire concentrator is providing flows for routing. A flow counter is maintained; when the number of active flows is 0, the softwire is deleted. Statistics are kept for both flows and softwires.

Softwire addresses are not specifically configured under any physical or virtual interface. Therefore, the number of established softwires does not affect throughput, and scalability is independent of the number of interfaces. The scalability is only limited to the number of flows that the platform (services DPC or PIC) can support.

This topic contains the following sections:

- [6to4 Overview on page 213](#)
- [DS-Lite Softwires—IPv4 over IPv6 on page 215](#)
- [6rd Softwires—IPv6 over IPv4 on page 216](#)

6to4 Overview

- [Basic 6to4 on page 214](#)
- [6to4 Anycast on page 214](#)
- [6to4 Provider-Managed Tunnels on page 215](#)

Basic 6to4

6to4 is an Internet transition mechanism for migrating from IPv4 to IPv6, a system that allows IPv6 packets to be transmitted over an IPv4 network (generally the IPv4 Internet) without the need to configure explicit tunnels. 6to4 is described in RFC 3056, *Connection of IPv6 Domains via IPv4 Clouds*. 6to4 is especially relevant during the initial phases of deployment to full, native IPv6 connectivity, since IPv6 is not required on nodes between the host and the destination. However, it is intended only as a transition mechanism and is not meant to be used permanently.

6to4 can be used by an individual host, or by a local IPv6 network. When used by a host, it must have a global IPv4 address connected, and the host is responsible for the encapsulation of outgoing IPv6 packets and the decapsulation of incoming 6to4 packets. If the host is configured to forward packets for other clients, often a local network, it is then a router.

There are two kinds of 6to4 virtual routers: border routers and relay routers. A 6to4 border router is an IPv6 router supporting a 6to4 pseudointerface, and is normally the border router between an IPv6 site and a wide-area IPv4 network. A relay router is a 6to4 router configured to support transit routing between 6to4 addresses and pure native IPv6 addresses.

In order for a 6to4 host to communicate with the native IPv6 Internet, its IPv6 default gateway must be set to a 6to4 address which contains the IPv4 address of a 6to4 relay router. To avoid the need for users to set this up manually, the Anycast address of 192.88.99.1 has been allocated to send packets to a 6to4 relay router. Note that when wrapped in 6to4 with the subnet and hosts fields set to zero, this IPv4 address (192.88.99.1) becomes the IPv6 address 2002:c058:6301::. To ensure BGP routing propagation, a short prefix of 192.88.99.0/24 has been allocated for routes pointed at 6to4 relay routers that use this Anycast IP address. Providers willing to provide 6to4 service to their clients or peers should advertise the Anycast prefix like any other IP prefix, and route the prefix to their 6to4 relay.

Packets from the IPv6 Internet to 6to4 systems must be sent to a 6to4 relay router by normal IPv6 routing methods. The specification states that such relay routers must only advertise 2002::/16 and not subdivisions of it to prevent IPv4 routes from polluting the routing tables of IPv6 routers. From there they can then be sent over the IPv4 Internet to the destination.

6to4 Anycast

Router 6to4 assumes that 6to4 routers and relays are managed and configured cooperatively. In particular, 6to4 sites must configure a relay router to carry the outbound traffic, which becomes the default IPv6 router (except for 2002::/16). The objective of the Anycast variant, defined in RFC 3068, *An Anycast Prefix for 6to4 Relay Routers*, is to avoid the need for such configuration. This makes the solution available for small or domestic users, even those with a single host or simple home gateway instead of a border router. This is achieved by defining 192.88.99.1 as the default IPv4 address for a 6to4 relay, and 2002:c058:6301:: as the default IPv6 router prefix (“well-known prefix”) for a 6to4 site.

RFC 6343, *Advisory Guidelines for 6to4 Deployment*, published in August 2011, identifies a wide range of problems associated with the use of unmanaged 6to4 Anycast relay routers.

6to4 Provider-Managed Tunnels

A solution to many problems associated with unmanaged Anycast 6to4 is presented in IETF informational draft draft-kuarsingh-v6ops-6to4-provider-managed-tunnel-02, *6to4 Provider-Managed Tunnels (PMT)*. That document, a “work in progress,” proposes a solution that allows providers to exercise greater control over the routing of 6to4 traffic.

Anycast 6to4 implies a default configuration for the user site. It does not require any particular user action. It does require an IPv4 Anycast route to be in place to a relay at 192.88.99.1. Traffic does not necessarily return to the same 6to4 gateway because of the the “well-known” 6to4 prefix used and advertised by all 6to4 traffic.

6to4 provider-managed tunnels (PMTs) facilitate the management of 6to4 tunnels using an Anycast configuration. 6to4 PMT enables service providers to improve 6to4 operation when network conditions provide suboptimal performance or break normal 6to4 operation. 6to4 PMT provides a stable provider prefix and forwarding environment by utilizing existing 6to4 relays with an added function of IPv6 prefix translation that controls the flow of return traffic.

The 6to4 managed tunnel model behaves like a standard 6to4 service between the customer IPv6 host or gateway and the 6to4 PMT relay (within the provider domain). The 6to4-PMT relay shares properties with 6rd (RFC5969) by decapsulating and forwarding embedded IPv6 flows, within an IPv4 packet, to the IPv6 Internet. The model provides an additional function which translates the source 6to4 prefix to a provider assigned prefix which is not found in 6rd (RFC5969) or traditional 6to4 operation. The 6to4-PMT relay provides a stateless (or stateful) mapping of the 6to4 prefix to a provider-supplied prefix by mapping the embedded IPv4 address in the 6to4 prefix to the provider prefix.

DS-Lite Softwires—IPv4 over IPv6

When an Internet service provider (ISP) begins to allocate new subscriber home IPv6 addresses and IPv6-capable equipment, dual-stack lite (DS-Lite) provides a method for the private IPv4 addresses behind the IPv6 customer edge (CE) WAN equipment to reach the IPv4 network. DS-Lite enables IPv4 customers to continue to access the Internet using their current hardware by using a softwire initiator, referred to as a Basic Bridging Broadband (B4), at the customer edge to encapsulate IPv4 packets into IPv6 packets and tunnel them over an IPv6 network to a softwire concentrator, referred to as an Address Family Transition Router (AFTR), for decapsulation. DS-Lite creates the IPv6 softwires that terminate on the services PIC. Packets coming out of the softwire can then have other services such as NAT applied on them.

DS-Lite is supported on Multiservices 100, 400, and 500 PICs on M Series routers and on MX Series routers equipped with Multiservices Dense Port Concentrator (DPCs).



NOTE: IPv6 Provider Edge (6PE), or MPLS-enabled IPv6, is available for ISPs with MPLS-enabled networks. These networks now can use Multiprotocol BGP (MP-BGP) to provide connectivity between the DS-Lite B4 and AFTR (or any two IPv6 nodes). DS-Lite properly handles encapsulation and decapsulation despite the presence of additional MPLS header information.

For more information on DS-Lite softwires, see the IETF draft *Dual Stack Lite Broadband Deployments Following IPv4 Exhaustion*.



NOTE: The most recent IETF draft documentation for DS-Lite uses new terminology:


- The term *software initiator* has been replaced by *B4*.
- The term *software concentrator* has been replaced by *AFTR*.

The Junos OS documentation generally uses the original terms when discussing configuration in order to be consistent with the command-line interface (CLI) statements used to configure DS-Lite.

6rd Softwires—IPv6 over IPv4

6rd software flow is shown in [Figure 15 on page 216](#).

Figure 15: 6rd Software Flow



The Junos OS supports a 6rd software concentrator on a services DPC or PIC to facilitate rapid deployment of IPv6 service to subscribers on native IPv4 CE WANs. IPv6 packets are encapsulated in IPv4 packets by a software initiator at the CE WAN. These packets are tunneled to a software concentrator residing on a multiservices DPC (branch relay). A software is created when IPv4 packets containing IPv6 destination information are received at the software concentrator, which decapsulates IPv6 packets and forwards them for IPv6 routing. All of these functions are performed in a single pass of the services PIC.

In the reverse path, IPv6 packets are sent to the Services DPC where they are encapsulated in IPv4 packets corresponding to the proper software and sent to the CE WAN.

The software concentrator creates softwires as the IPv4 packets are received from the CE WAN side or IPv6 packets are received from the Internet. A 6rd software on the Services DPC is identified by the 3-tuple containing the service set ID, CE software initiator IPv4 address, and software concentrator IPv4 address. IPv6 flows are also created for the encapsulated IPv6 payload, and are associated with the specific software that carried

them in the first place. When the last IPv6 flow associated with a softwire ends, the softwire is deleted. This simplifies configuration and there is no need to create or manage tunnel interfaces.

6rd is supported on Multiservices 100, 400, and 500 PICs on M Series and T Series routers, and on MX Series platforms equipped with Multiservices DPCs.

For more information on 6rd softwires, see RFC 5969, *IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification*.

**Related
Documentation**

- *Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards*
- [Configuring a 6rd Softwire Concentrator on page 245](#)
- [Configuring a DS-Lite Softwire Concentrator on page 227](#)
- [Configuring Softwire Rules on page 219](#)
- [Configuring Service Sets for Softwire on page 220](#)

Softwires Configuration Overview

- [Configuring Software Rules on page 219](#)
- [Configuring Service Sets for Software on page 220](#)

Configuring Software Rules

You configure software rules to instruct the router how to direct traffic to the addresses specified for 6rd or DS-Lite software concentrators. Software rules do not perform any filtration of the traffic. They do not include a **from** statement, and the only option in the **then** statement is to specify the address of the 6rd or DS-Lite software concentrator.

You can create a software rule consisting of one or more terms and associate a particular 6rd or DS-Lite software concentrator with each term. You can include the software rule in service sets along with other services rules.

To configure a software rule:

1. Assign a name to the rule.

```
[edit services software ]  
user@host# edit rule rule-name
```

2. Specify the match direction.

```
[edit services software rule rule-name]  
user@host# set match-direction (input | output)
```

3. Assign a name for the first term.

```
[edit services software rule rule-name]  
user@host# edit term term-name
```

4. Associate a 6rd or DS-Lite software concentrator with this term.

```
[edit services software rule rule-name term term-name]  
user@host# set then ds-lite name
```

Or

```
user@host# set then v6rd v6rd-software-concentrator
```

5. Repeat Steps 3 and 4 for as many additional terms as needed.

Related Documentation

- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 213](#)
- [Configuring a 6rd Software Concentrator on page 245](#)
- [Configuring a DS-Lite Software Concentrator on page 227](#)
- [Configuring IPv6 Multicast Interfaces](#)
- [Configuring Service Sets for Software on page 220](#)

Configuring Service Sets for Software

You must include software rules or a software rule set in a service set to enable software processing. You must include a stateful firewall rule for DS-Lite.

To configure service sets for software:

1. Include a software rule or rule set in the service set.

```
[edit services service-set service-set-name]  
user@host# set software-rules rule software-rule-name
```

2. When using a 6rd software, include a stateful-firewall rule.

```
[edit services service-set service-set-name]  
user@host# set stateful-firewall-rules software-rule-name
```

3. You can include a NAT rule for flows originated by DS-Lite softwares.



NOTE:

Currently a NAT rule configuration is required with a DS-Lite software configuration when you use interface service set configurations; NAT is not required when using next-hop service set configurations. NAT processing from IPv4 to IPv6 address pools and vice versa is not currently supported. FTP, HTTP, and RSTP are supported.



NOTE: With a DS-Lite software concentrator, if you configure stateful firewall rules without configuring NAT rules, using an interface service set causes the ICMP echo reply messages to be not sent correctly to DS-Lite. This behavior occurs if you apply a service set to both inet and inet6 families. In such a scenario, the traffic that is not destined to the DS-Lite software concentrator is also processed by the service set and the packets might be dropped, although the service set must not process such packets.

To prevent the problem to incorrect processing of traffic applicable for DS-Lite, you must configure a next-hop style service set and not an interface style service set. This problem does not occur when you configure NAT rules with interface service sets for DS-Lite.

For further information, see ““[Configuring Service Rules](#)” on page 36.”

**Related
Documentation**

- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 213](#)
- [Configuring Softwire Rules on page 219](#)
- [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 234](#)

Transitioning to IPv6 Using 6to4 Softwires

- [Configuring a 6to4 Provider-Managed Tunnel on page 223](#)

Configuring a 6to4 Provider-Managed Tunnel

When configuring a 6to4 provider-managed tunnel (PMT), replace the Anycast destination with the address of a managed relay in the provider network.

To configure a 6to4 PMT:

1. Configure the ingress interface for 6to4 traffic. Include the name of the service set that identifies the rules for input and output service on this interface.

```
[edit interfaces ge-0/2/1]
user@host# set unit logical-unit-number family family service input service-set-name
user@host# set unit logical-unit-number family family service output service-set-name
user@host# set unit logical-unit-number family family address address
```

For example:

```
[edit interfaces ge-0/2/1]
user@host# set unit 0 family inet service input service-set v6to4-pmt
user@host# set unit 0 family inet service output service-set v6to4-pmt
user@host# set unit 0 family inet address 130.130.130.1/24
```

2. Configure the egress interface.

```
[edit interfaces ge-0/2/2]
user@host# set unit logical-unit-number family family address address
```

For example:

```
[edit interfaces ge-0/2/2]
user@host# set unit 0 family inet6 address 4ABC::1/16
```

3. Configure the service interface that contains the rules for processing incoming traffic. Include a syslog option and associate a logical unit.

```
[edit interfaces sp-2/0/0]
user@host# edit services-options syslog host host-name services any
user@host# edit unit logical-unit-number family family
user@host# edit unit 0 family family
```

For example:

```
[edit interfaces sp-2/0/0]
```

```
user@host# set services-options syslog host local services any
user@host# set unit 0 family inet
user@host# set unit 0 family inet6
```

4. Configure the softwire concentrator and softwire rule for 6to4. In the Junos OS, 6to4 PMT configuration uses the same options as 6rd.

```
[edit services softwire softwire-concentrator v6rd v6to4]
user@host# set softwire-address softwire-address
user@host# set ipv4-prefix ipv4-prefix
user@host# set v6rd-prefix v6rd-prefix
user@host# set mtu-v4 mtu-v4
```

For example:

```
[edit services softwire softwire-concentrator v6rd v6to4]
user@host# set softwire-address 192.88.99.1
user@host# set ipv4-prefix 130.130.130.2/32
user@host# set v6rd-prefix 2002::0/16
user@host# set mtu-v4 9192
```

5. Define the softwire rule that will process traffic on the ingress interface.

```
[edit services softwire rule v6to4-r1]
user@host# set match-direction input
user@host# set term term-name then v6rd softwire-concentrator
```

For example:

```
[edit services softwire rule v6to4-r1]
user@host# set match-direction input
user@host# set term t1 then v6rd v6to4
```

6. Define a stateful firewall rule that will accept all incoming traffic on the ingress interface.

```
[edit services stateful-firewall rule sfw-r1]
user@host# set match-direction direction
user@host# set term term-name then accept
user@host# set term term-name then syslog
```

For example:

```
[edit services stateful-firewall rule sfw-r1]
user@host# set match-direction input-output
user@host# set term t1 then accept
user@host# set term t1 then syslog
```

7. Define the NAT pool to be used for IPv6 NAT translation. This pool supports translation of the Anycast 6to4 relay addresses to addresses at the provider-managed relay.

```
[edit services nat pool v6to4-pmt]
user@host# set address address
user@host# port automatic
```

For example:

```
[edit services nat pool v6to4-pmt]
user@host# set address 3ABC::1/128
user@host# set port automatic
```

8. Define the NAT rule for translation.

```
[edit services nat rule rule-name]  
user@host# set match-direction input  
user@host# set term term-name then translated source-pool pool-name  
user@host# set term t1 then translated translation-type translation-type
```

For example:

```
[edit services nat rule v6to4-pmt-r1]  
user@host# set match-direction input  
user@host# set term t1 then translated source-pool v6to4-pmt  
user@host# set term t1 then translated translation-type napt-66
```

9. Define the service set that specifies the softwire rule and NAT rule.

```
[edit services service-set v6to4-pmt]  
user@host# set softwire-rules rule-name  
user@host# set stateful-firewall-rules rule-name  
user@host# set nat-rules rule-name  
user@host# set interface-service service-interface interface-name
```

For example:

```
[edit services service-set v6to4-pmt]  
user@host# set softwire-rules v6to4-r1  
user@host# set stateful-firewall-rules sfw-r1  
user@host# set nat-rules v6to4-pmt-r1  
user@host# set interface-service service-interface sp-2/0/0
```


Transitioning to IPv6 Using DS-Lite Softwires

- [Configuring a DS-Lite Softwire Concentrator on page 227](#)
- [Example: Basic DS-Lite Configuration on page 228](#)
- [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 234](#)
- [Protecting CGN Devices Against Denial of Service \(DOS\) Attacks on page 241](#)
- [DS-Lite Subnet Limitation on page 242](#)

Configuring a DS-Lite Softwire Concentrator

To configure a DS-Lite softwire concentrator:

1. Assign a name to the DS-Lite softwire concentrator.

```
[edit services softwire softwire-concentrator]
user@host# edit ds-lite ds-lite-softwire-concentrator
```

2. Specify the address of the softwire tunnel.

```
[edit services softwire softwire-concentrator ds-lite ds-lite-softwire-concentrator]
user@host# set softwire-address address
```

3. Specify the MTU for the softwire tunnel.

```
[edit services softwire softwire-concentrator ds-lite ds-lite-softwire-concentrator]
user@host# set mtu-v6 mtu-v6
```



NOTE: This option sets the maximum transmission unit when encapsulating IPv4 packets into IPv6. If the final length is greater than the MTU, the IPv6 packet will be fragmented. This option is mandatory since it depends on other network parameters under administrator control.

4. To copy DSCP information from the IPv6 header into the decapsulated IPv4 header, include the **copy-dscp** statement.

```
[edit services softwire softwire-concentrator ds-lite ds-lite-softwire-concentrator]
user@host# set copy-dscp
```

5. Specify the maximum number of flows for the softwire

```
[edit services software-concentrator ds-lite ds-lite-software-concentrator]  
user@host# set flow-limit 1000
```

**Related
Documentation**

- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 213](#)
- [Configuring Software Rules on page 219](#)
- [Configuring IPv6 Multicast Interfaces](#)
- [Configuring Service Sets for Software on page 220](#)
- [Example: Basic DS-Lite Configuration on page 228](#)
- [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 234](#)

Example: Basic DS-Lite Configuration

- [Requirements on page 228](#)
- [Configuration Overview and Topology on page 228](#)
- [Configuration on page 229](#)

Requirements


The following hardware components can perform DS-Lite:

- M Series Multiservice Edge routers with Multiservices PICs
- T Series Core routers with Multiservices PICs
- MX Series 3D Universal Edge routers with Multiservices DPCs

Configuration Overview and Topology

This example describes how to configure an MX Series router with an MS-DPC as an AFTR to facilitate the flow shown in [Figure 16 on page 229](#).

Figure 16: DS-Lite Topology



In this example, the DS-Lite softwire concentrator, or AFTR, is an MX Series router with two Gigabit interfaces and a Services DPC. The interface facing the B4 element is ge-3/1/5 and the interface facing the Internet is ge-3/1/0.

Configuration

- [Chassis Configuration on page 229](#)
- [Interfaces Configuration on page 229](#)
- [Network Address and Port Translation Configuration on page 231](#)
- [Softwire Configuration on page 232](#)
- [Service Set Configuration on page 233](#)

Chassis Configuration

Step-by-Step Procedure

To configure the service PIC (FPC 0 Slot 0) with the Layer 3 service package:

1. Enter the **edit chassis** hierarchy level.

```
user@host# edit chassis
```
2. Configure the Layer 3 service package.

```
[edit chassis]
user@host# set fpc 0 pic 0 adaptive-services service-package layer-3
```

Interfaces Configuration

Step-by-Step Procedure

To configure interfaces facing the B4 (softwire initiator) and facing the Internet:

1. Go the **[edit interfaces]** edit hierarchy level for ge-3/1/0, which faces the Internet.

```
host# edit interfaces ge-3/1/0
```

2. Define the interface.

```
[edit interfaces ge-3/1/0]
user@host# set description AFTR-Internet
user@host# set unit 0 family inet address 128.0.0.2/24
```

3. Go to the **[edit interfaces]** hierarchy level for ge-3/1/5, which faces the B4.

```
user@host# up 1
[edit]
user@host# edit interfaces ge-3/1/5
```

4. Define the interface.

```
[edit interfaces ge-3/1/5]
user@host# set description AFTR-B4
user@host# set unit 0 family inet
user@host# edit unit 0 family inet6
[edit unit 0 family inet6]
user@host# set service input service-set sset
user@host# set service output service-set sset
user@host# set address 2001:0:0:2::1/48
```

5. Go to the **[edit interfaces]** hierarchy level for sp-0/0/0, used to host the DS-Lite AFTR.

```
[edit]
user@host# edit interfaces sp-0/0/0
```

6. Define the interface.

```
[edit interfaces sp-0/0/0]
user@host# set description AFTR-B4
user@host# set unit 0 family inet
user@host# edit unit 0 family inet6
```



```

Results user@host# show interfaces ge-3/1/0
description AFTR-Internet;
unit 0 {
    family inet {
        address 128.0.0.2/24;
    }
}

user@host# show interfaces ge-3/1/5
description AFTR-B4;
unit 0 {
    family inet;
    family inet6 {
        service {
            input {
                service-set sset;
            }
            output {
                service-set sset;
            }
        }
        address 2001:0:0:2::1/48;
    }
}

user@host# show interfaces sp-o/o/o
unit 0 {
    family inet;
    family inet6;
}

```

Network Address and Port Translation Configuration

Step-by-Step Procedure

To configure NAPT:

1. Go to the **[edit services nat]** hierarchy level.

```

user@host# edit services nat
[edit services nat]

```
2. Define a NAT pool p1.

```

user@host# set pool p1 address 129.0.0.1/32 port automatic

```
3. Define a NAT rule, beginning with the match direction.

```

[edit services nat]
user@host# set rule r1 match-direction input

```
4. Define a **term** for the rule, beginning with a **from** clause.

```

[edit services nat]
user@host# set rule r1 term t1 from source-address 10.0.0.0/16

```
5. Define the desired translation in a **then** clause. In this case, use dynamic source translation.

```

[edit services nat]
user@host# set rule r1 term t1 then translated source-pool p1 translation-type napt-44

```
6. (Optional) Configure logging of translation information for the rule.

```
[edit services nat]
user@host# set rule r1 term t1 then syslog
```

```
Results user@host# show services nat
pool p1 {
  address 129.0.0.1/32;
  port {
    automatic;
  }
}
rule r1 {
  match-direction input;
  term t1 {
    from {
      source-address {
        10.0.0.0/16;
      }
    }
    then {
      translated {
        source-pool p1;
        translation-type {
          napt-44;
        }
      }
      syslog;
    }
  }
}
```

Software Configuration

Step-by-Step Procedure To configure the DS-Lite software concentrator and associated rules:

1. Go to the **[edit services software]** hierarchy level.

```
user@host# edit services software
```
2. Define the DS-Lite software concentrator.

```
[edit services software]
user@host# set software-concentrator ds-lite ds-1 software-address 1001::1 mtu-v6 1460
```
3. Define the software rule.

```
[edit services software]
user@host# set rule r1 match-direction input term t1 then ds-lite ds1.
```

Results

```

user@host# show services software
software-concentrator {
  ds-lite ds1 {
    software-address 1001::1;
    mtu-v6 1460;
  }
}
rule r1 {
  match-direction input;
  term t1 {
    then {
      ds-lite ds1;
    }
  }
}

```

Service Set Configuration

Step-by-Step Procedure Configure a service set that includes software and NAT rules and specifies either interface-service or next-hop service. This example uses a next-hop service.

1. Go to the **[edit services service-set]** hierarchy level, naming the service set.

```
user@host# edit services service-set sset
```

2. Define the NAT rule to be used for IPv4-to-IPv4 translation.

```

[edit services service-set sset]
user@host# set nat-rules r1

```

3. Define the software rule to define the software tunnel.

```

[edit services service-set sset]
user@host# set software-rules r1

```

4. Define the interface service,

```

[edit services service-set sset]
user@host# set interface-service service-interface sp-0/0/0.0

```



TIP: In order to avoid or minimize IPv6 fragmentation, you can configure a TCP maximum segment size (MSS) for your service set.

5. (Optional) Define a TCP MSS.

```

[edit services service-set sset]
user@host# set tcp-mss 1024

```

Results `user@host# show services service-set`

```
syslog {
  host local {
    services any;
  }
}
software-rules r1;
nat-rules r1;
interface-service {
  service-interface sp-0/0/0;
}
```

- Related Documentation**
- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 213](#)
 - [Configuring a DS-Lite Software Concentrator on page 227](#)
 - [Configuring Software Rules on page 219](#)
 - [Configuring Service Sets for Software on page 220](#)
 - [Example: Basic 6rd Configuration on page 247](#)
 - [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 234](#)

Example: Configuring DS-Lite and 6rd in the Same Service Set

- [Requirements on page 234](#)
- [Overview on page 234](#)
- [Configuration on page 234](#)

Requirements

The following hardware components can perform DS-Lite:

- M Series Multiservice Edge routers with Multiservices PICs
- T Series Core routers with Multiservices PICs
- MX Series 3D Universal Edge routers with Multiservices DPCs

Overview

This example describes a software solution that includes DS-Lite and 6rd in the same service set.

Configuration

Chassis Configuration

Step-by-Step Procedure

To configure the chassis:

1. Configure the ingress interface.

```
user@host# edit interfaces ge-1/2/0
[edit interfaces ge-1/2/0]
```

```

user@host# set unit 0 family inet service input service-set v6rd-dslite-service-set
user@host# set unit 0 family inet service output service-set v6rd-dslite-service-set
user@host# set unit 0 family inet address address 10.10.10.1/24
user@host# set unit 0 family inet6 service input service-set v6rd-dslite-service-set
user@host# set unit 0 family inet6 service output service-set v6rd-dslite-service-set
user@host# set unit 0 family inet6 address address address 2001::1/16

```

Here the service set is applied on the inet (IPv4) and inet6 (IPv6) families of subunit 0. Both DS-Lite IPv6 traffic and 6rd IPv4 traffic hits the service filter and is sent to the services PIC.

2. Configure the egress interface (IPv6 Internet). The IPv4 server that the DS-Lite clients are trying to reach is at 200.200.200.2/24, and the IPv6 server is at 3ABC::2/16.

```

user@host# edit interfaces ge-1/2/2
[edit interfaces ge-1/2/2]
user@host# set unit 0 family inet address 200.200.200.1/24
user@host# set unit 0 family inet6 address 3ABC::1/16

```

3. Configure the services PIC.

```

user@host# edit interfaces sp-3/0/0
[edit interfaces sp-3/0/0]
user@host# set unit 0 family inet
user@host# set unit 0 family inet6

```

Results [edit interfaces]
user@host# show
ge-1/2/0 {
 unit 0 {
 family inet {
 service {
 input {
 service-set v6rd-dslite-service-set;
 }
 output {
 service-set v6rd-dslite-service-set;
 }
 }
 address 10.10.10.1/24;
 }
 family inet6 {
 service {
 input {
 service-set v6rd-dslite-service-set;
 }
 output {
 service-set v6rd-dslite-service-set;
 }
 }
 address 2001::1/16;
 }
 }
}
ge-1/2/2 {
 unit 0 {
 family inet {
 address 200.200.200.1/24;
 }
 family inet6 {
 address 3ABC::1/16;
 }
 }
}
sp-3/0/0 {
 unit 0 {
 family inet;
 family inet6;
 }
}

Software Concentrator, Software Rule, Stateful Firewall Rule Configuration

Step-by-Step Procedure

To configure the software concentrator, software rule, and stateful firewall rule:

1. Configure the DS-Lite and 6rd software concentrators.

```
user@host# edit services software software-concentrator ds-lite ds1
[edit services software software-concentrator ds-lite ds1]
user@host# set software-address 1001::1
user@host# mtu-v6 9192
user@host# up 1
user@host# edit v6rd v6rd-dom1
[edit services software software-concentrator v6rd v6rd-dom1]
user@host# set software-address 30.30.30.1
```

```

user@host# set ipv4-prefix 10.10.10.0/24
user@host# set v6rd-prefix 3040::0/16
user@host# set mtu-v4 9192

```

2. Configure the softwire rules.

```

user@host# edit services softwire rule v6rd-r1]
[edit services softwire rule v6rd-r1]
user@host# set match-direction input
user@host# set term t1 then v6rd v6rd-dom1
user@host# up 1
user@host# edit services softwire]
[edit services softwire]
user@host# edit rule dslite-r1
[edit services softwire rule dslite-r1]
user@host# set term dslite-t1 then ds-lite ds1

```

The following routes are added by the services PIC daemon on the Routing Engine:

```

user@router# run show route 30.30.30.1
inet.0: 43 destinations, 46 routes (42 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

30.30.30.1/32      *[Static/786432] 00:24:11
                  Service to v6rd-dslite-service-set

```

```
[edit]
```

```
user@router# run show route 3040::0/16
```

```

inet6.0: 23 destinations, 33 routes (23 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

3040::/16          *[Static/786432] 00:24:39
                  Service to v6rd-dslite-service-set

```

```

user@router# run show route 1001::1
inet6.0: 33 destinations, 43 routes (33 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

1001::1/128       *[Static/1] 1w2d 22:05:41
                  Service to v6rd-dslite-service-set

```

3. Configure a stateful firewall rule.

```

user@host# edit services stateful-firewall rule r1
[edit services stateful-firewall rule r1]
user@host# set match-direction input-output
user@host# set term t1 then accept

[edit services stateful-firewall]
rule r1 {
  match-direction input-output;
  term t1 {
    then {
      accept;
    }
  }
}

```

Results [edit services software]
user@host# **show**
software-concentrator {
 ds-lite ds1 {
 software-address 1001::1;
 mtu-v6 9192;
 }
 v6rd v6rd-dom1 {
 software-address 30.30.30.1;
 ipv4-prefix 10.10.10.0/24;
 v6rd-prefix 3040::0/16;
 mtu-v4 9192;
 }
}
rule v6rd-r1 {
 match-direction input;
 term t1 {
 then {
 v6rd v6rd-dom1;
 }
 }
}
rule dslite-r1 {
 match-direction input;
 term dslite-t1 {
 then {
 ds-lite ds1;
 }
 }
}
}

[edit services stateful-firewall]
user@host# **show**
rule r1 {
 match-direction input-output;
 term t1 {
 then {
 accept;
 }
 }
}

NAT Configuration for DS-Lite

Step-by-Step Procedure

To configure NAT for DS-Lite:

1. Configure a NAT pool for DS-Lite.

```
user@host# edit services nat pool dslite-pool  
[edit services nat pool dslite-pool]  
user@host# set address-range low 33.33.33.1 high 33.33.33.32  
user@host# set port automatic
```
2. Configure a NAT rule.

```
user@host# up 1  
[edit services nat rule dslite-nat-r1]  
user@host# set match-direction input
```



```
user@host# set term dslite-nat-t1 from source-address 20.20.0.0/16 then translated  
translation-type napt-44
```

Results [edit services nat]
user@host# show
pool dslite-pool {
 address-range low 33.33.33.1 high 33.33.33.32;
 port {
 automatic;
 }
}
rule dslite-nat-r1 {
 match-direction input;
 term dslite-nat-t1 {
 from {
 source-address {
 20.20.0.0/16;
 }
 }
 then {
 translated {
 source-pool dslite-pool;
 translation-type {
 source dynamic;
 }
 }
 }
 }
}

Because of this NAT rule, the following NAT routes are installed for the reverse DS-Lite traffic:

```
user@router# run show route 33.33.33.0/24
inet.0: 48 destinations, 52 routes (47 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

33.33.33.1/32      *[Static/1] 1w2d 23:08:38
                  Service to v6rd-dslite-service-set
33.33.33.2/31      *[Static/1] 1w2d 23:08:38
                  Service to v6rd-dslite-service-set
33.33.33.4/30      *[Static/1] 1w2d 23:08:38
                  Service to v6rd-dslite-service-set
33.33.33.8/29      *[Static/1] 1w2d 23:08:38
                  Service to v6rd-dslite-service-set
33.33.33.16/28     *[Static/1] 1w2d 23:08:38
                  Service to v6rd-dslite-service-set
33.33.33.32/32     *[Static/1] 1w2d 23:08:38
                  Service to v6rd-dslite-service-set
```

The NAT rule triggers address translation for the traffic coming from 20.20.0.0/16 to public address range 33.33.33.1 to 33.33.33.32.

Service Set Configuration

Step-by-Step Procedure

This service set has a stateful firewall rule and 6rd rule for 6rd service. The service set also includes a softwire rule for DS-Lite and a NAT rule to perform address translation for all DS-Lite traffic. The NAT rule performs NAPT translation in the forward direction on the source address and port of the DS-Lite traffic.

To configure the service set:

1. Define the service set.

```
user@host# edit services service-set v6rd-dslite-service-set
```
2. Configure the service set rules.

```
[edit services service-set v6rd-dslite-service-set]  
user@host# set software-rules dslite-r1  
user@host# set stateful-firewall-rules r1  
user@host# set nat-rules dslite-nat-r1
```
3. Configure the service set interface-service.

```
[edit services service-set v6rd-dslite-service-set]  
user@host# set interface-service service-interface sp-3/0/0
```

Results

```
[edit services service-set]  
user@host# show  
v6rd-dslite-service-set {  
  software-rules v6rd-r1;  
  software-rules dslite-r1;  
  stateful-firewall-rules r1;  
  nat-rules dslite-nat-r1;  
  interface-service {  
    service-interface sp-3/0/0;  
  }  
}
```

- Related Documentation**
- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 213](#)
 - [Configuring Service Sets for Softwire on page 220](#)
 - [Example: Basic DS-Lite Configuration on page 228](#)
 - [Example: Basic 6rd Configuration on page 247](#)

Protecting CGN Devices Against Denial of Service (DOS) Attacks

You can now choose configuration options that help prevent or minimize the effect of attempted denial of service (DOS) attacks.

- [Mapping Refresh Behavior on page 241](#)
- [EIF Inbound Flow Limit on page 242](#)

Mapping Refresh Behavior

Prior to the implementation of the new options for configuring NAT mapping refresh behavior, described in this topic, a conversation was kept alive when either inbound or outbound flows were active. This remains the default behavior. You can now also specify mapping refresh for only inbound flows or only outbound flows. To configure mapping refresh behavior, include the **mapping-refresh (inbound | outbound | inbound-outbound)** statement at the **[edit services nat rule *rule-name* term *term-name* then translated secure-nat-mapping]** hierarchy level.

EIF Inbound Flow Limit

Previously, the number of inbound connections on an EIF mapping was limited only by the maximum flows allowed on the system. You can now configure the number of inbound flows allowed for an EIF. To limit the number of inbound connections on an EIF mapping, include the **eif-flow-limit *number-of-flows*** statement at the **[edit services nat rule *rule-name* term *term-name* then translated secure-nat-mapping]** hierarchy level.

DS-Lite Subnet Limitation

- [DS-Lite Per Subnet Limitation Overview on page 242](#)
- [Configuring DS-Lite Per Subnet Session Limitation to Prevent Denial of Service Attacks on page 243](#)

DS-Lite Per Subnet Limitation Overview

Junos OS enables you to limit the number of software flows from a subscriber's basic bridging broadband (B4) device at a given point in time, preventing subscribers from excessive use of addresses within the subnet. This limitation reduces the risk of denial-of-service (DoS) attacks.

A household using IPv6 with DS-Lite is a subnet, not just an individual IP address. The subnet limitation feature associates a subscriber and mapping with an IPv6 prefix instead of an IPv6 address. A subscriber can use any IPv6 addresses in that prefix as a DS-Lite B4 address and potentially exhaust carrier-grade NAT resources. The subnet limitation feature enables greater control of resource utilization by identifying a subscriber with a prefix instead of a specific address.

The subnet limit provides the following features:

- Flows utilize the complete B4 address.
- Prefix length can be configured per service set under software-options for the individual service-set.
- Port blocks are allocated per prefix of the subscriber B4 device, and not on each B4 address (if the prefix length is less than 128). If prefix the length is 128, then each IPv6 address is treated as a B4. Port blocks are allocated per 128-bit V6 address.
- Session limit, defined under the DSLite software concentrator configuration, limits the number of IPv4 sessions for the prefix.
- EIM, EIF, and PCP mappings are created per software tunnel (full 128 bit IPv6 address). Stale mappings time out based on timeout values.
- If prefix length is configured, then PCP **max-mappings-per-subscriber** (configurable under **pcp-server**) is based on the prefix only, and not the full B4 address.
- SYSLOGS for PBA allocation and release contain the prefix portion of the address completed with all zeros. SYSLOGS for PCP alloc and release, Flow creation and deletion will still contain the complete IPv6 address.

The **show services nat mappings address-pooling-paired** operational command output now shows the mapping for the prefix. The mapping shows the address of the active B4.

The **show services softwire statistics ds-lite** output includes a new field that displays the number of times the session limit was exceeded for the MPC.

Configuring DS-Lite Per Subnet Session Limitation to Prevent Denial of Service Attacks

To configure DS-Lite per subnet session limitation:

1. Configure the size of the subnet prefix to which limiting is applied. Specify a prefix length of 56, 64, 96, or 128.

```
[edit]
user@router# set services service-set service-set-name softwire-options
                        dslite-ipv6-prefix-length 56.
```



NOTE: Ensure that all mappings are cleared before changing the prefix length.

2. Configure the maximum number of subscriber sessions allowed per prefix. You can configure from 0 through 16,384 sessions.

```
[edit]
user@router# set services softwire softwire-concentrator dslite
                        dslite-concentrator-name session-limit-per-prefix 12
```



NOTE: You cannot use **flow-limit** and **session-limit-per-prefix** in the same **dslite** configuration.

Transitioning to IPv6 Using 6rd Softwires

- [Configuring a 6rd Software Concentrator on page 245](#)
- [Configuring Stateful Firewall Rules for 6rd Software on page 246](#)
- [Example: Basic 6rd Configuration on page 247](#)
- [Inter-Chassis High Availability for MS-MIC and MS-MPC on page 252](#)
- [High Availability and Load Balancing for 6rd Softwires on page 264](#)

Configuring a 6rd Software Concentrator

To configure a 6rd software concentrator:

1. Assign a name to the 6rd software concentrator.

```
[edit services software software-concentrator]  
user@host# edit v6rd v6rd-software-concentrator
```

2. Specify the address of the software tunnel.

```
[edit services software software-concentrator v6rd v6rd-software-concentrator]  
user@host# set software-address address
```

3. Specify the MTU for the software tunnel.

```
[edit services software software-concentrator v6rd v6rd-software-concentrator]  
user@host# set mtu-v4 mtu-v4
```



TIP: In this release there is no support for fragmentation and reassembly, therefore the MTUs on the IPv6 and IPV4 network must be properly configured by the administrator.



NOTE: Configuration changes to 6rd software concentrators do not become effective in the Packet Forwarding Engine. This is a known limitation. If you attempt to add the new configuration of software concentrators by overriding the existing configuration of 1024 software concentrators, which is the maximum limit of software concentrators that the system supports, the new configuration is not updated. To work around this limitation, you must delete the existing configuration and commit the settings, and then add the new configuration of software concentrators and commit the settings.



NOTE: For 6rd software concentrators, packet drops are observed and error messages logged on the virtual terminal session (VTY) console, if one inline services (si-) interface is replaced with another si- interface without stopping the traffic during the replacement of the interface. In a scenario in which an si- interface is associated with a service set that has a large number of software concentrators, replacing that interface without halting the traffic causes traffic disruption. You must stop the traffic and restart it during such a replacement of si- interfaces with 6rd software concentrators. The following error messages are displayed on the VTY console of the FPC:

packet discarded because no ifl or not SI ifl

Related Documentation

- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 213](#)
- [Configuring Software Rules on page 219](#)
- [Configuring Stateful Firewall Rules for 6rd Software on page 246](#)
- [Configuring Service Sets for Software on page 220](#)
- [Example: Basic 6rd Configuration on page 247](#)
- [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 234](#)

Configuring Stateful Firewall Rules for 6rd Software

You must configure a stateful firewall rule for use with 6rd softwires. The stateful firewall service is used only to direct packets to the software, not for firewalling purposes. The 6rd software service itself must be stateless. To support stateless processing, you must include an **allow** term in both directions of the stateful firewall policy.

To include a stateful firewall rule for 6rd software processing:

1. Assign a name to the rule.

```
[edit services stateful-firewall]  
user@host# edit rule rule-name
```

2. Specify the match direction.

```
[edit services stateful-firewall rule-name]
```



```
user@host# set match-direction input-output
```

3. Assign a name for the term.

```
[edit services stateful-firewall rule-name]
```

```
user@host# edit term term-name
```

4. Specify that all traffic in both directions should be accepted for the software process.

```
[edit services stateful-firewall rule-name term term-name]
```

```
user@host# set then accept
```

Related Documentation

- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 213](#)
- [Configuring a 6rd Software Concentrator on page 245](#)
- [Configuring Software Rules on page 219](#)
- [Configuring Service Sets for Software on page 220](#)
- [Example: Basic 6rd Configuration on page 247](#)
- [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 234](#)

Example: Basic 6rd Configuration

- [Requirements on page 247](#)
- [Overview on page 247](#)
- [Configuration on page 247](#)

Requirements

This example describes how a 6rd concentrator can be configured for a 6rd domain, D1, to provide IPv6 Internet connectivity.

The following hardware components can perform 6rd:

- M Series Multiservice Edge routers with Multiservices PICs
- T Series Core routers with Multiservices PICs
- MX Series 3D Universal Edge routers with Multiservices DPCs

Overview

This configuration example describes how to configure a basic 6rd tunneling solution.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces ge-1/2/0 unit 0 family inet service input service-set v6rd-dom1-service-set
set interfaces ge-1/2/0 unit 0 family inet service output service-set v6rd-dom1-service-set
```

```
set interfaces ge-1/2/0 unit 0 family inet address 10.10.10.1/24
set interfaces ge-1/2/0 unit 0 family inet6 service input service-set v6rd-dom1-service-set
set interfaces ge-1/2/0 unit 0 family inet6 service output service-set v6rd-dom1-service-set
set interfaces ge-1/2/2 unit 0 family inet6 address 3abc::1/16
set interfaces sp-0/2/0 unit 0 family inet
set interfaces sp-0/2/0 unit 0 family inet6
set services software software-concentrator v6rd v6rd-dom1 software-address 30.30.30.1
set services software software-concentrator v6rd v6rd-dom1 ipv4-prefix 10.10.10.0/24
set services software software-concentrator v6rd v6rd-dom1 v6rd-prefix 3040::0/16
set services software software-concentrator v6rd v6rd-dom1 mtu-v4 9192
set services software rule v6rd-dom1 match-direction input
set services software rule v6rd-dom1 term t1 then v6rd v6rd-dom1
set services service-set v6rd-dom1-service-set software-rules v6rd-dom1
set services service-set v6rd-dom1-service-set stateful-firewall-rules r1
set services service-set v6rd-dom1-service-set interface-service service-interface sp-0/2/0
set services stateful-firewall rule r1 match-direction input-output
set services stateful-firewall rule r1 term t1 then accept
```

Chassis Configuration

Step-by-Step Procedure

To configure the chassis:

1. Define the ingress interface.

```
user@host# edit interfaces ge-1/2/0
```
2. Configure the ingress interface logical unit and input/output service options.

```
[edit interfaces ge-1/2/0]
user@host# set unit 0 family inet service input service-set v6rd-dom1-service-set
user@host# set unit 0 family inet service output service-set v6rd-dom1-service-set
user@host# set unit 0 family inet6 service input service-set v6rd-dom1-service-set
user@host# set unit 0 family inet6 service output service-set v6rd-dom1-service-set
```
3. Configure the address of the ingress interface.

```
[edit interfaces ge-1/2/0]
user@host# set unit 0 family inet address 10.10.10.1/24
```
4. Define the egress interface.

```
user@host# up
[edit interfaces]
user@host# edit ge-1/2/2
```
5. Define the logical unit and address for the egress interface.

```
[edit interfaces ge-1/2/2]
user@host# set unit 0 family inet6 address 3ABC::1/16
```
6. Define the services PIC.

```
[edit interfaces ge-1/2/2]
user@host# up
[edit interfaces]
user@host# edit sp-0/2/0
```
7. Configure the logical unit for the services PIC.

```
[edit interfaces sp-0/2/0]
```

```

user@host# up
[edit interfaces]
user@host# set unit 0 family inet
user@host# set unit 0 family inet6

```

Results

```

[edit interfaces]
user@router# show
sp-0/2/0 {
    unit 0 {
        family inet;
        family inet6;
    }
}
ge-1/2/0 {
    unit 0 {
        family inet {
            service {
                input {
                    service-set v6rd-dom1-service-set;
                }
                output {
                    service-set v6rd-dom1-service-set;
                }
            }
            address 10.10.10.1/24;
        }
        family inet6 {
            service {
                input {
                    service-set v6rd-dom1-service-set;
                }
                output {
                    service-set v6rd-dom1-service-set;
                }
            }
        }
    }
}
ge-1/2/2 {
    unit 0 {
        family inet6 {
            address 3abc::1/16;
        }
    }
}

```

Software Concentrator, Software Rule, and Stateful Firewall Rule Configuration

Step-by-Step Procedure

To configure the software concentrator, software rule, and stateful firewall rule:

1. Define the 6rd software concentrator.

```

user@host# top
user@host# edit services software software-concentrator v6rd v6rd-dom1

```

2. Configure the software concentrator properties. Here, software address 30.30.30.1 is the software concentrator IPv4 address, 10.10.10.0/24 is the IPv4 prefix of the CE WAN side, and 3040::0/16 is the IPv6 prefix of the 6rd domain D1.

```
[edit services software software-concentrator v6rd v6rd-dom1]
user@host# set software-address 30.30.30.1
user@host# set ipv4-prefix 10.10.10.0/24
user@host# set v6rd-prefix 3040::0/16
user@host# set mtu-v4 9192
```

3. Define the software rule.

```
[edit services software software-concentrator v6rd v6rd-dom1]
user@host# up 2
[edit services software]
user@host# edit rule v6rd-dom1
[edit services software rule v6rd-dom1]
user@host# set match-direction input
[edit services software rule v6rd-dom1]
user@host# set term t1 then v6rd v6rd-dom1
```

4. Define a stateful firewall rule and properties. You must configure a stateful firewall rule that accepts all traffic in both the input and output direction in order for 6rd to work; however, this is not enforced through the CLI. This is because in IPv6, gratuitous IPv6 packets are expected (due to Anycast) and should not be dropped. The service PIC can handle reverse traffic without seeing all forward traffic. This can also happen with service PIC switchover in the middle of a session. By default, the stateful firewall on the service PIC will drop all traffic unless a rule is configured explicitly to allow it.

```
[edit services software software-concentrator v6rd v6rd-dom1]
user@host# up 3
[edit services]
user@host# edit services stateful-firewall
[edit services stateful-firewall]
user@host# edit rule r1
[edit services stateful-firewall rule r1]
user@host# set match-direction input-output
user@host# set term t1 then accept
```

Results [edit services software]
 user@router# show
 software-concentrator {
 v6rd v6rd-dom1 {
 software-address 30.30.30.1;
 ipv4-prefix 10.10.10.0/24;
 v6rd-prefix 3040::0/16;
 mtu-v4 9192;
 }
 }
 rule v6rd-dom1-r1 {
 match-direction input;
 term t1 {
 then {
 v6rd v6rd-dom1;
 }
 }
 }

Service Set Configuration

Step-by-Step Procedure To configure the service set:

1. Define the service set for 6rd processing.

```
user@host# top
user@host# edit services service-set v6rd-dom1-service-set
```
2. Define the software and stateful firewall rules for the service set.

```
[edit services service-set v6rd-dom1-service-set]
user@host# set software-rules v6rd-dom1
user@host# set stateful-firewall-rules r1
```
3. Define the interface-service for the service set.

```
[edit services service-set v6rd-dom1-service-set]
user@host# set interface-service service-interface sp-0/2/0
```

Results [edit service-set v6rd-dom1-service-set]
 user@host# show
 software-rules v6rd-dom1-r1
 interface-service {
 service-interface sp-0/2/0;
 }

- Related Documentation**
- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 213](#)
 - [Configuring a 6rd Software Concentrator on page 245](#)
 - [Configuring Software Rules on page 219](#)
 - [Configuring Stateful Firewall Rules for 6rd Software on page 246](#)
 - [Configuring Service Sets for Software on page 220](#)
 - [Example: Basic DS-Lite Configuration on page 228](#)
 - [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 234](#)

Inter-Chassis High Availability for MS-MIC and MS-MPC

Inter-chassis high availability supports stateful synchronization of services using a switchover to a backup services PIC on a different chassis. The feature is described in the following topics:

- [Inter-Chassis High Availability for Stateful Firewall and NAPT44 Overview \(MS-MIC, MS-MPC\) on page 252](#)
- [Configuring Inter-Chassis High Availability for Stateful Firewall and NAPT44 \(MS-MPC, MS-MIC\) on page 253](#)
- [Example: Inter-Chassis Stateful High Availability for NAT and Stateful Firewall \(MS-MIC, MS-MPC\) on page 254](#)

Inter-Chassis High Availability for Stateful Firewall and NAPT44 Overview (MS-MIC, MS-MPC)

Carrier-grade NAT (CGN) deployments can use dual-chassis implementations to provide a redundant data path and redundancy for key components in the router. Although intra-chassis high availability can be used in dual-chassis environments, it deals only with service PIC failures. If traffic is switched to a backup router due to some other failure in the router, state is lost. Inter-chassis high availability preserves state and provides redundancy using fewer service PICs than intra-chassis high availability. Only long-lived flows are synchronized between the master and backup chassis in the high availability pair. The service PICs do not replicate state until an explicit CLI command, **request services redundancy (synchronize | no-synchronize)**, is issued to start or stop the state replication. Stateful firewall, NAPT44, and APP state information can be synchronized.



NOTE: When both the master and backup PICs are up, replication starts immediately when the **request services redundancy** command is issued.

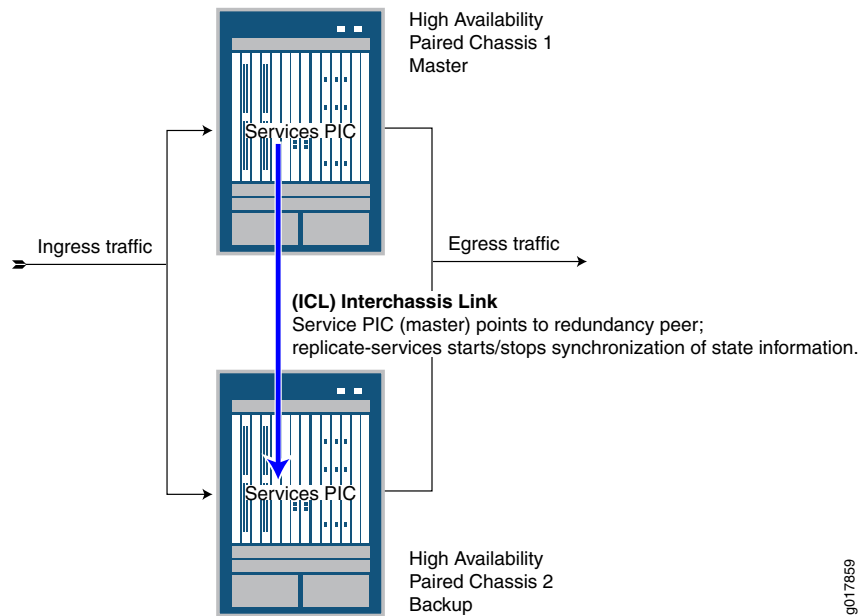
In order to use Inter-chassis high availability, you must use service sets configured for next-hop service interfaces. Inter-chassis high availability works with ms- service interfaces configured on MS-MIC or MS-MPC interface cards. A unit other than unit 0 must be configured with the **ip-address-owner service-plane** option.

The following restrictions apply:

- NAPT44 is the only translation type supported.
- Checkpointing is not supported for ALGs, PBA port block allocation (PBA), endpoint-independent mapping (EIM), or endpoint-independent filters (EIF).

[Figure 17 on page 253](#) shows the inter-chassis high availability topology.

Figure 17: Inter-Chassis High Availability Topology



Configuring Inter-Chassis High Availability for Stateful Firewall and NAPT44 (MS-MPC, MS-MIC)

To configure inter-chassis availability for stateful firewall and NAPT44 on MS-MIC or MS-MPC service PICs, perform the following configuration steps on each chassis of the high availability pair:

1. At the `[edit interfaces interface-name redundancy-options]` hierarchy level, set the **ipaddress** for the **redundancy-peer**. This IPv4 address specifies one of the hosted IP addresses of the remote PIC. This address is used by the TCP channel between the HA pairs.

```
[edit interfaces interface-name redundancy-options]
user@host# set redundancy-peer ipaddress ipaddress
```



NOTE: When you enable or disable high availability of MS-MICs or MS-MPCs by configuring or removing the primary and backup adaptive services PICs by using the `redundancy-options redundancy-peer ipaddress address` statement at the `[edit interfaces interface-name]` hierarchy level, the configuration change is treated as a catastrophic event for each service-set that refers to the affected interface at the `[edit services service-set name interface-service service-interface interface-name]` hierarchy level. A catastrophic event at the service-set level has the effect of deactivating the service set, applying the change, and then reactivating the service set.

2. Specify the name of a special routing instance, or VRF, you want applied to the HA synchronization traffic between the high availability pair.

```
[edit interfaces interface-name redundancy-options]
```

```
user@host# set routing-instance instance-name
```

3. For the service set defining an interface that is a member of the high availability pair, configure the service replication options using the *replicate-services* option.

```
[edit services service-set service-set-name replicate-services]  
user@host# set replication-threshold threshold-value  
stateful-firewall  
nat
```

Example: Inter-Chassis Stateful High Availability for NAT and Stateful Firewall (MS-MIC, MS-MPC)

This example shows how to configure inter-chassis high availability for stateful firewall and NAT services.

- [Requirements on page 254](#)
- [Overview on page 254](#)
- [Configuration on page 255](#)

Requirements

This example uses the following hardware and software components:

- Two MX480 routers with MS-MPC line cards
- Junos OS Release 13.3 or later

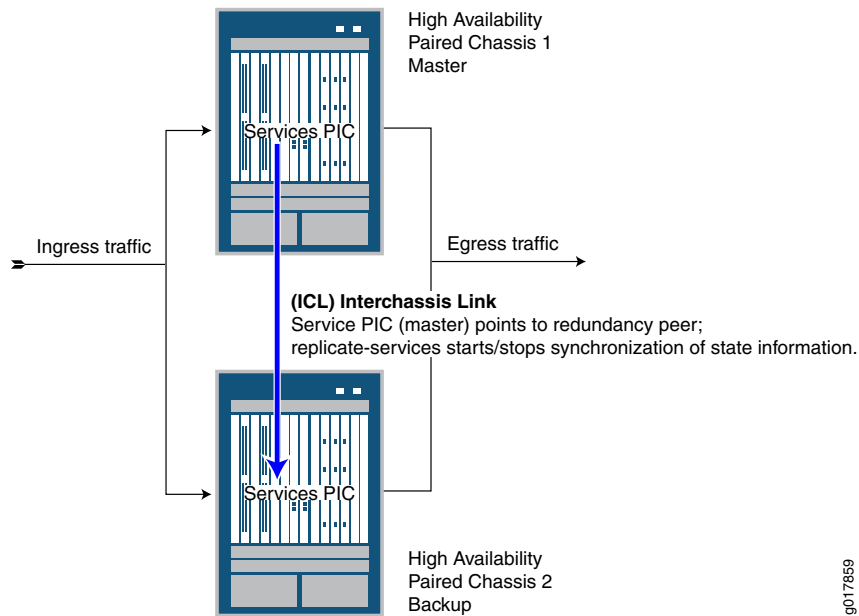
Overview

Two MX 3D routers are identically configured to facilitate stateful failover for firewall and NAT services in case of a chassis failure.

Topology

[Figure 18 on page 255](#) shows the inter-chassis high availability topology.

Figure 18: Inter-Chassis High Availability Topology



g017859

Configuration

To configure inter-chassis high availability for this example, perform these tasks:

- [Configuring Interfaces for Chassis 1 on page 257](#)
- [Configure Routing Information for Chassis 1 on page 258](#)
- [Configuring NAT and Stateful Firewall for Chassis 1 on page 259](#)
- [Configuring the Service Set on page 260](#)
- [Configuring Interfaces for Chassis 2 on page 261](#)
- [Configure Routing Information for Chassis 2 on page 263](#)

CLI Quick Configuration

To quickly configure this example on the routers, copy the following commands and paste them into the router terminal window after removing line breaks and substituting interface information specific to your site.



NOTE: The following configuration is for chassis 1.

```
[edit]
set interfaces ms-4/0/0 redundancy-options redundancy-peer ipaddress 5.5.5.2
set interfaces ms-4/0/0 redundancy-options routing-instance HA
set interfaces ms-4/0/0 unit 10 ip-address-owner service-plane
set interfaces ms-4/0/0 unit 10 family inet address 5.5.5.1/32
set interfaces ms-4/0/0 unit 20 family inet
set interfaces ms-4/0/0 unit 20 service-domain inside
set interfaces ms-4/0/0 unit 30 family inet
set interfaces ms-4/0/0 unit 30 service-domain outside
set interfaces ge-2/0/0 vlan-tagging
```

```

set interfaces ge-2/0/0 unit 0 vlan-id 100 family inet address 20.1.1.1/24
set routing-instances HA instance-type vrf
set routing-instances HA interface ge-2/0/0.0
set routing-instances HA interface ms-4/0/0.10
set routing-instances HA route-distinguisher 1:1
set policy-options policy-statement dummy term 1 then reject
set routing-instances HA vrf-import dummy
set routing-instances HA vrf-export dummy
set routing-instances HA routing-options static route route 5.5.1/32 next-hop
  ms-4/0/0.10
set routing-instances HA routing-options static route route 5.5.2/32 next-hop 20.1.1.2
set services nat pool p2 address 32.0.0.0/24
set services nat pool p2 port automatic random-allocation
set services nat pool p2 address-allocation round-robin
set services nat rule r2 match-direction input
set services nat rule r2 term t1 from source-address 129.0.0.0/8
set services nat rule r2 term t1 from source-address 128.0.0.0/8
set services nat rule r2 term t1 then translated source-pool p2
set services nat rule r2 term t1 then translated translation-type napt-44
set services nat rule r2 term t1 then translated address-pooling paired
set services nat rule r2 term t1 then syslog
set services stateful-firewall rule r2 match-direction input
set services stateful-firewall rule r2 term t1 from source-address any-unicast
set services stateful-firewall rule r2 term t1 then accept
set services stateful-firewall rule r2 term t1 then syslog
set services service-set ss2 replicate-services replication-threshold 180
set services service-set ss2 replicate-services stateful-firewall
set services service-set ss2 replicate-services nat
set services service-set ss2 stateful-firewall-rules r2
set services service-set ss2 nat-rules r2
set services service-set ss2 next-hop-service inside-service-interface ms-4/0/0.20
set services service-set ss2 next-hop-service outside-service-interface ms-4/0/0.30
set services service-set ss2 syslog host local class session-logs
set services service-set ss2 syslog host local class stateful-firewall-logs
set services service-set ss2 syslog host local class nat-logs

```



NOTE: The following configuration is for chassis 2. The NAT, stateful firewall, and service-set information must be identical for chassis 1 and 2.

```

set interfaces ms-4/0/0 redundancy-options routing-instance HA
set interfaces ms-4/0/0 unit 10 ip-address-owner service-plane
set interfaces ms-4/0/0 unit 10 family inet address 5.5.2/32
set interfaces ms-4/0/0 unit 20 family inet
set interfaces ms-4/0/0 unit 20 service-domain inside
set interfaces ms-4/0/0 unit 30 family inet
set interfaces ms-4/0/0 unit 30 service-domain outside
set interfaces ge-2/0/0 vlan-tagging
set interfaces ge-2/0/0 unit 0 vlan-id 100 family inet address 20.1.1.2/24
set routing-instances HA instance-type vrf
set routing-instances HA interface ge-2/0/0.0
set routing-instances HA interface ms-4/0/0.10
set routing-instances HA route-distinguisher 1:1
set routing-instances HA vrf-import dummy

```

```

set routing-instances HA vrf-export dummy
set routing-instances HA routing-options static route 5.5.5.2/32 next-hop ms-4/0/0.10
set routing-instances HA routing-options static route 5.5.5.1/32 next-hop 20.1.1.1
set services nat pool p2 address 32.0.0.0/24
set services nat pool p2 port automatic random-allocation
set services nat pool p2 address-allocation round-robin
set services nat rule r2 match-direction input
set services nat rule r2 term t1 from source-address 129.0.0.0/8
set services nat rule r2 term t1 from source-address 128.0.0.0/8
set services nat rule r2 term t1 then translated source-pool p2
set services nat rule r2 term t1 then translated translation-type napt-44
set services nat rule r2 term t1 then translated address-pooling paired
set services nat rule r2 term t1 then syslog
set services stateful-firewall rule r2 match-direction input
set services stateful-firewall rule r2 term t1 from source-address any-unicast
set services stateful-firewall rule r2 term t1 then accept
set services stateful-firewall rule r2 term t1 then syslog
set services service-set ss2 replicate-services replication-threshold 180
set services service-set ss2 replicate-services stateful-firewall
set services service-set ss2 replicate-services nat
set services service-set ss2 stateful-firewall-rules r2
set services service-set ss2 nat-rules r2
set services service-set ss2 next-hop-service inside-service-interface ms-4/0/0.20
set services service-set ss2 next-hop-service outside-service-interface ms-4/0/0.30
set services service-set ss2 syslog host local class session-logs
set services service-set ss2 syslog host local class stateful-firewall-logs
set services service-set ss2 syslog host local class nat-logs

```

Configuring Interfaces for Chassis 1.

Step-by-Step Procedure

The interfaces for each of the HA pair of routers are configured identically with the exception of the following service PIC options:

- `redundancy-options redundancy-peer ipaddress address`
- `unit unit-number family inet address address` of a unit, other than 0, that contains the `ip-address-owner service-plane` option

To configure interfaces:

1. Configure the redundant service PIC on chassis 1.

```

[edit interfaces]
user@host# set interfaces ms-4/0/0 redundancy-options redundancy-peer
ipaddress 5.5.5.2
user@host# set interfaces ms-4/0/0 redundancy-options routing-instance HA
user@host# set interfaces ms-4/0/0 unit 10 ip-address-owner service-plane
user@host# set interfaces ms-4/0/0 unit 10 family inet address 5.5.5.1/32
user@host# set interfaces ms-4/0/0 unit 20 family inet
user@host# set interfaces ms-4/0/0 unit 20 service-domain inside
user@host# set interfaces ms-4/0/0 unit 30 family inet
user@host# set interfaces ms-4/0/0 unit 30 service-domain outside

```

2. Configure the interfaces for chassis 1 that are used as interchassis links for synchronization traffic.

```

user@host# set interfaces ge-2/0/0 vlan-tagging

```

```
user@host# set interfaces ge-2/0/0 unit 0 vlan-id 100 family inet address 20.1.1.1/24
```

3. Configure remaining interfaces as needed.

Results

```
user@host# show interfaces
ge-2/0/0 {
    vlan-tagging;
    unit 0 {
        vlan-id 100;
        family inet {
            address 20.1.1.1/24;
        }
    }
}
ms-4/0/0 {
    redundancy-options {
        redundancy-peer {
            ipaddress 5.5.5.2;
        }
        routing-instance HA;
    }
    unit 10 {
        ip-address-owner service-plane;
        family inet {
            address 5.5.5.1/32;
        }
    }
    unit 20 {
        family inet;
        family inet6;
        service-domain inside;
    }
    unit 30 {
        family inet;
        family inet6;
        service-domain outside;
    }
}
```

Configure Routing Information for Chassis 1

Step-by-Step Procedure Detailed routing configuration is not included for this example. A routing instance is required for the HA synchronization traffic between the chassis as follows:

- Configure routing instances for Chassis 1.


```
user@host# set routing-instances HA instance-type vrf
user@host# set routing-instances HA interface ge-2/0/0.0
user@host# set routing-instances HA interface ms-4/0/0.10
user@host# set routing-instances HA route-distinguisher 1:1
user@host# set policy-options policy-statement dummy term 1 then reject
user@host# set routing-instances HA vrf-import dummy
user@host# set routing-instances HA vrf-export dummy
@user@host# set routing-instances HA routing-options static route route 5.5.5.1/32
next-hop ms-4/0/0.10
user@host# set routing-instances HA routing-options static route route 5.5.5.2/32
next-hop 20.1.1.2
```

```

Results  @user@host# show routing-instances
HA {
    instance-type vrf;
    interface ge-2/0/0.0;
    interface ms-4/0/0.10;
    route-distinguisher 1:1;
    vrf-import dummy;
    vrf-export dummy;
    routing-options {
        static {
            route 5.5.5.1/32 next-hop ms-4/0/0.10;
            route 5.5.5.2/32 next-hop 20.1.1.2;
        }
    }
}

```

Configuring NAT and Stateful Firewall for Chassis 1

Step-by-Step Procedure Configure NAT and stateful firewall identically on both routers. To configure NAT and stateful firewall:

1. Configure NAT as needed.

```

user@host# set services nat pool p2 address 32.0.0.0/24
user@host# set services nat pool p2 port automatic random-allocation
user@host# set services nat pool p2 address-allocation round-robin
user@host# set services nat rule r2 match-direction input
user@host# set services nat rule r2 term t1 from source-address 129.0.0.0/8
user@host# set services nat rule r2 term t1 from source-address 128.0.0.0/8
user@host# set services nat rule r2 term t1 then translated source-pool p2
user@host# set services nat rule r2 term t1 then translated translation-type napt-44
user@host# set services nat rule r2 term t1 then translated address-pooling paired
user@host# set services nat rule r2 term t1 then syslog

```

2. Configure stateful firewall as needed.

```

user@host# set services stateful-firewall rule r2 match-direction input
user@host# set services stateful-firewall rule r2 term t1 from source-address
any-unicast
user@host# set services stateful-firewall rule r2 term t1 then accept
user@host# set services stateful-firewall rule r2 term t1 then syslog

```

```
Results user@host# show services nat
nat {
    pool p2 {
        address 32.0.0.0/24;
        port {
            automatic {
                random-allocation;
            }
        }
        address-allocation round-robin;
    }
    rule r2 {
        match-direction input;
        term t1 {
            from {
                source-address {
                    129.0.0.0/8;
                    128.0.0.0/8;
                }
            }
            then {
                translated {
                    source-pool p2;
                    translation-type {
                        napt-44;
                    }
                    address-pooling paired;
                }
                syslog;
            }
        }
    }
}

user@host# show services stateful-firewall
rule r2 {
    match-direction input;
    term t1 {
        from {
            source-address {
                any-unicast;
            }
        }
        then {
            accept;
            syslog;
        }
    }
}
```

Configuring the Service Set

Step-by-Step Procedure Configure the the service set identically on both routers. To configure the service set:

1. Configure the service set replication options.

```
user@host# set services service-set ss2 replicate-services replication-threshold
180
```

```
user@host# set services service-set ss2 replicate-services stateful-firewall
user@host# set services service-set ss2 replicate-services nat
```

2. Configure references to NAT and stateful firewall rules for the service set.

```
user@host# set services service-set ss2 stateful-firewall-rules r2
user@host# set services service-set ss2 nat-rules r2
```

3. Configure next-hop service interface on the MS-PIC.

```
user@host# set services service-set ss2 next-hop-service inside-service-interface
ms-4/0/0.20
user@host# set services service-set ss2 next-hop-service outside-service-interface
ms-4/0/0.30
```

4. Configure desired logging options.

```
user@host# set services service-set ss2 syslog host local class session-logs
user@host# set services service-set ss2 syslog host local class stateful-firewall-logs
user@host# set services service-set ss2 syslog host local class nat-logs
```

Results

```
user@host# show services service-set ss2
syslog {
    host local {
        class {
            session-logs;
            inactive: stateful-firewall-logs;
            nat-logs;
        }
    }
}
replicate-services {
    replication-threshold 180;
    stateful-firewall;
    nat;
}
stateful-firewall-rules r2;
inactive: nat-rules r2;
next-hop-service {
    inside-service-interface ms-3/0/0.20;
    outside-service-interface ms-3/0/0.30;
}
}
```

Configuring Interfaces for Chassis 2

Step-by-Step Procedure The interfaces for each of the HA pair of routers are configured identically with the exception of the following service PIC options:

- **redundancy-options redundancy-peer ipaddress address**
- **unit unit-number family inet address address** of a unit, other than 0, that contains the **ip-address-owner service-plane** option

1. Configure the redundant service PIC on chassis 2.

The **redundancy-peer ipaddress** points to the address of the unit (unit 10) on ms-4/0/0 on chassis 1 that contains the **ip-address-owner service-plane** statement.

```
[edit interfaces]
set interfaces ms-4/0/0 redundancy-options redundancy-peer ipaddress 5.5.5.1
user@host# set interfaces ms-4/0/0 redundancy-options routing-instance HA
user@host# set interfaces ms-4/0/0 unit 10 ip-address-owner service-plane
user@host# set interfaces ms-4/0/0 unit 10 family inet address 5.5.5.2/32
user@host# set interfaces ms-4/0/0 unit 20 family inet
user@host# set interfaces ms-4/0/0 unit 20 service-domain inside
user@host# set interfaces ms-4/0/0 unit 30 family inet
user@host# set interfaces ms-4/0/0 unit 30 service-domain outside
```

2. Configure the interfaces for chassis 2 that are used as interchassis links for synchronization traffic

```
user@host# set interfaces ge-2/0/0 vlan-tagging
user@host# set interfaces ge-2/0/0 unit 0 vlan-id 100 family inet address 20.1.1.2/24
```

3. Configure remaining interfaces for chassis 2 as needed.

Results

```

user@host# show interfaces
ms-4/0/0 {
    redundancy-options {
        redundancy-peer {
            ipaddress 5.5.5.1;
        }
        routing-instance HA;
    }
    unit 0 {
        family inet;
    }
    unit 10 {
        ip-address-owner service-plane;
        family inet {
            address 5.5.5.2/32;
        }
    }
}
ge-2/0/0 {
    vlan-tagging;
    unit 0 {
        vlan-id 100;
        family inet {
            address 20.1.1.2/24;
        }
    }
    unit 10 {
        vlan-id 10;
        family inet {
            address 2.10.1.2/24;
        }
    }
}

```

Configure Routing Information for Chassis 2

Step-by-Step Procedure Detailed routing configuration is not included for this example. A routing instance is required for the HA synchronization traffic between the two chassis and is included here.

- Configure routing instances for chassis 2.


```

user@host# set routing-instances HA instance-type vrf
user@host# set routing-instances HA interface ge-2/0/0.0
user@host# set routing-instances HA interface ms-4/0/0.10
user@host# set routing-instances HA route-distinguisher 1:1
user@host# set policy-options policy-statement dummy term 1 then reject
user@host# set routing-instances HA vrf-import dummy
user@host# set routing-instances HA vrf-export dummy
user@host# set routing-instances HA routing-options static route 5.5.5.2/32 next-hop ms-4/0/0.10
user@host# set routing-instances HA routing-options static route 5.5.5.1/32 next-hop 20.1.1.1

```



NOTE: The following configuration steps are *identical* to the steps shown for chassis 1.

- Configuring NAT and Stateful Firewall
- Configuring the Service Set

Results @user@host# **show services routing-instances**

```
HA {
    instance-type vrf;
    interface xe-2/2/0.0;
    interface ms-4/0/0.10;
    route-distinguisher 1:1;
    vrf-import dummy;
    vrf-export dummy;
    routing-options {
        static {
            route 5.5.5.2/32 next-hop ms-4/0/0.10;
            route 5.5.5.1/32 next-hop 20.1.1.1;
        }
    }
}
```

High Availability and Load Balancing for 6rd Softwires

- [Load Balancing a 6rd Domain Across Multiple Services PICs on page 264](#)
- [Example: Load Balancing a 6rd Domain Across Multiple Services PICs on page 264](#)
- [Configuring High Availability for 6rd Using 6rd Anycast on page 269](#)

Load Balancing a 6rd Domain Across Multiple Services PICs

The 6rd domain is an IPv6 network, which can potentially be very large. A single PIC, or network processing unit (NPU) on a Multiservices DPC, might not be able to handle all the traffic for the 6rd domain. To alleviate load problems, you can load-balance the 6rd domain traffic across multiple PICs. To do so, assign the same software rule to different services sets that use different interfaces. Configure explicit routes and equal-cost multipath (ECMP) to load-balance the 6rd traffic.

Example: Load Balancing a 6rd Domain Across Multiple Services PICs

- [Hardware and Software Requirements on page 264](#)
- [Overview on page 265](#)
- [Configuration on page 265](#)

Hardware and Software Requirements

This example requires the following hardware:

- An MX Series 3D Universal Edge router with a services DPC with two available NPUs or an M Series Multiservice Edge router with two services PICs available for 6rd software concentrator processing
- A domain name server (DNS)

This example uses the following software:

- Junos OS Release 11.4 or higher

Overview

Because of anticipated volume, a provider needs to balance 6rd software traffic between two services PICs.

Configuration

- [Chassis Configuration on page 265](#)
- [Software Concentrator and Software Rule Configuration on page 266](#)
- [Stateful Firewall Configuration on page 266](#)
- [Service Set Configuration on page 267](#)
- [Load-Balancing Configuration on page 267](#)

Chassis Configuration

Step-by-Step Procedure

To configure the chassis:

1. Define the ingress interface and its properties.


```
user@host# edit interfaces ge-1/2/0
user@host# set unit 0 family inet address 10.10.10.1/16
```
2. Define the egress interface and its properties. In this example, the IPv6 clients try to reach the IPv6 server at 3abc::2/16.


```
user@host# edit interfaces ge-1/2/2
user@host# set unit 0 family inet6 address 3ABC::1/16
```
3. Define the services PICs for selection as software concentrators by the load-balancing process. This configuration uses two PICs/NPUs: sp-3/0/0 and sp-3/1/0. A next-hop style service set is configured (shown in the next section).


```
user@host# edit interfaces sp-3/0/0
[edit interfaces ge-3/0/0]
user@host# set services-options syslog host local services any
user@host# set unit 0 family inet
user@host# set unit 0 family inet6
user@host# set unit 1 family inet service-domain inside
user@host# set unit 1 family inet service-domain outside
user@host# set unit 2 family inet service-domain inside
user@host# set unit 2 family inet service-domain outside
user@host# up 1
[edit]
user@host# edit interfaces sp-3/1/0
[edit interfaces sp-3/1/0]
user@host# set services-options syslog host local services any
user@host# set unit 0 family inet
user@host# set unit 0 family inet6
user@host# set unit 1 family inet service-domain inside
user@host# set unit 1 family inet service-domain outside
user@host# set unit 2 family inet service-domain inside
user@host# set unit 2 family inet service-domain outside
```

Software Concentrator and Software Rule Configuration

Step-by-Step Procedure The software configuration is straightforward. In this example, the 6rd domain prefix is 3040::0/16, the 6rd software concentrator IPv4 address is 30.30.30.1, and the customer IPv4 network is 10.10.0.0/16. In the customer premises equipment (CPE) network, all customer edge (CE) devices have addresses that belong to the 10.10.0.0/16 network. To configure the software:

1. Go to the **[edit services software]** hierarchy level.
`user@host# edit services software`
2. Configure IPv6 multicast.
`[edit services software]
user@host# set ipv6-multicast-interfaces all`
3. Go to the software concentrator v6rd hierarchy level and name the software concentrator **shenick01-rd1**.
`[edit services software]
user@host# edit software-concentrator v6rd shenick01-rd1`
4. Configure the software concentrator properties.
`[edit services software software-concentrator v6rdshenick01-rd1]
user@host# set software-address 30.30.30.1
user@host# set ipv4-prefix 10.10.0.0/16
user@host# set v6rd-prefix 3040::/16
user@host# set mtu-v4 9192`
5. Configure a software rule for incoming 6rd traffic.
`[edit services software software-concentrator v6rd shenick01-rd1]
user@host# up 1
[edit services software]
user@host# edit rule shenick01-r1
[edit services software rule shenick01-r1]
user@host# set match-direction input
user@host# set term t1 then v6rd shenick01-rd1`

Stateful Firewall Configuration

Step-by-Step Procedure To configure the stateful firewall rule:

1. Go to the stateful firewall hierarchy level and define a rule.
`user@host# edit services stateful-firewall rule r1`
2. Set the match direction.
`[edit services stateful-firewall rule r1]
user@host# set match-direction input-output`
3. Configure a term that accepts all traffic.
`[edit services stateful-firewall rule r1]
user@host# set term t1 then accept`

Service Set Configuration

Step-by-Step Procedure This configuration provides two service sets, each pointing to a different network processing unit (NPU). Both service sets use the same stateful firewall and softwire rules. Because they use the same softwire rule, they refer to same 6rd softwire concentrator. This results in the software concentrator being hosted on both the NPUs.

To configure the service set:

1. Define a service set for the first NPU.

```
user@host# edit services service-set v6rd-sset1
```
2. Configure the softwire and stateful firewall rules for the first NPU.

```
[edit services service-set v6rd-sset1]
user@host# set softwire-rules shenick01-r1
user@host# set stateful-firewall-rules r1
```
3. Configure the inside and outside interfaces for the next-hop service.

```
[edit services service-set v6rd-sset1]
user@host# set next-hop-service inside-service-interface sp-3/0/0.1
user@host# set next-hop-service outside-service-interface sp-3/0/0.2
```
4. Define a service set for the second NPU.

```
user@host# edit services service-set v6rd-sset2
```
5. Configure the softwire and stateful firewall rules for the second NPU.

```
[edit services service-set v6rd-sset2]
user@host# set softwire-rules shenick01-r1
user@host# set stateful-firewall-rules r1
```
6. Configure the inside and outside interfaces for the next-hop service.

```
[edit services service-set v6rd-sset1]
user@host# set next-hop-service inside-service-interface sp-3/1/0.1
user@host# set next-hop-service outside-service-interface sp-3/1/0.2
```

Load-Balancing Configuration

Step-by-Step Procedure To configure load balancing:
 Configure explicit routes and ECMP to load-balance the 6rd traffic. Configure explicit routes for both the 6rd concentrator IPv4 address and the 6rd domain prefix, so that they point to both NPUs.

1. To configure static routes for the 6rd domain using the routing-table inet6.0, go to the **[edit forwarding-options rib inet6.0 static]** hierarchy level and set the routes for the 6rd domain and the 6rd concentrator IPv4 address.

```
user@host edit forwarding-options rib inet6.0 static
[edit forwarding-options rib inet6.0 static]
user@host# set route 3040::0/16 next-hop [ sp-3/0/0.2 sp-3/1/0.2 ]
user@host# set route 30.30.30.1/32 next-hop [ sp-3/0/0.1 sp-3/1/0.1 ]
```

The service PIC daemon (spd) also adds default routes to these addresses pointing to the NPUs. However, the routes added by the spd use different metrics, which are computed based on the FPC, PIC, slot numbers, and subunit of the services PIC if used in the service set configuration. The static routes configured in this sample configuration will have metrics of 5 and therefore a higher preference than the spd-added routes.

The explicitly configured routes are as follows:

```
root@router# run show route 30.30.30.1
inet.0: 37 destinations, 40 routes (36 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
30.30.30.1/32      *[Static/5] 00:00:10
                   > via sp-3/0/0.1
                   via sp-3/1/0.1
                   [Static/786433] 00:23:03
                   > via sp-3/0/0.1
                   [Static/851969] 00:00:09
                   > via sp-3/1/0.1
```

```
root@router# run show route 3040::/16
inet6.0: 20 destinations, 33 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
3040::/16         *[Static/5] 00:00:15
                   via sp-3/0/0.2
                   > via sp-3/1/0.2
                   [Static/786434] 00:23:08
                   > via sp-3/0/0.2
                   [Static/851970] 00:00:14
                   > via sp-3/1/0.2
```



BEST PRACTICE: The spd-installed routes have higher metric values (hence a low preference) and the metrics are different. If the metrics are different and ECMP is not enabled, even though multiple routes exist for the same destination, only one of the routes is picked up all the time (based on the metric). For ECMP you must configure equal-cost routes, and hence a manual configuration of routes is needed as shown above.

2. Configure equal-cost multipath (ECMP) load balancing by configuring the hash key at the **[edit forwarding-options hash-key]** hierarchy level.

```
user@host# forwarding-options hash-key
[edit forwarding-options hash-key]
user@host# set family inet layer-3 destination-address
user@host# set family inet layer-3 source-address
user@host# set family inet6 layer-3 destination-address
user@host# set family inet6 layer-3 source-address
```

3. Verify your configuration by displaying **forwarding-options**.

```
user@host# show forwarding-options
hash-key {
    family inet { <== IPv4 traffic from CEs uses this
```

```

        layer-3 {
            destination-address;
            source-address;
        }
    }
    family inet6 { <== IPv6 traffic from Internet uses this
        layer-3 {
            destination-address;
            source-address;
        }
    }
}

```



TIP: Both IPv4 and IPv6 hash keys must be configured. The IPv4 hash key is used to distribute the traffic coming from CPE devices to the 6rd branch relay. The IPv6 hash key is used to distribute the traffic coming from the IPv6 Internet to the 6rd domain. Because the hash in the forward and reverse direction is for different families, different flows from the same session can reside on different NPUs. However, 6rd processing is stateless (as far as mapping IPv6 packets to softwires is concerned), so this should not be a problem.

Configuring High Availability for 6rd Using 6rd Anycast

You configure 6rd Anycast by defining two service sets that use the same softwire rule in both service sets, just as you do when you configure load balancing for 6rd. However, you do not configure ECMP, and as a result, the services PIC daemon (spd) installs two routes *each* for the softwire concentrator address and 6rd domain pointing to each service interface. The forwarding plane can select any route based on the priority, which is computed when the spd installs the routes. The priority is computed based on the FPC, PIC, slot numbers, and subunit number used on the sp- interface. *Only one PIC is used* based on the route priority, and that PIC gets all of the 6rd traffic. If the PIC goes down, the route pointing to it is also deleted and the forwarding plane automatically selects the alternate available PIC.

6rd Anycast is completely stateless. The spd installs the route and doesn't run any state machine for the PIC. Because the routes are pre-installed and service sets are already on the PIC, there is no service delay if a failover occurs.

Related Documentation

- *Junos Address Aware Carrier Grade NAT and IPv6 Feature Guide*

Monitoring and Troubleshooting Softwires

- [Ping and Traceroute for DS-Lite on page 271](#)
- [Monitoring Softwire Statistics on page 271](#)
- [Monitoring CGN, Stateful Firewall, and Softwire Flows on page 273](#)

Ping and Traceroute for DS-Lite

With Junos OS Release 11.4, you can use the **ping** and **traceroute** commands to determine the status of the DS-Lite softwire tunnels:

- **IPv6 ping**—The softwire address endpoint on the DS-Lite softwire terminator (AFTR) is usually configured only at the **[edit services softwire]** hierarchy level; it need not be hosted on any interface. Previous releases of the Junos OS software did not provide replies to pings to the IPv6 softwire address when the AFTR was not configured on a specific interface or loopback. An IPv6 ping enables the softwire initiator (B4) to verify the softwire address of the AFTR before creating a tunnel.
- **IPv4 ping**—A special IPv4 address, 192.0.0.1, is reserved for the AFTR. Previous releases of the Junos OS did not respond to any pings sent to this address. A B4 and other IPv4 nodes can now ping to this address to determine whether the DS-Lite tunnel is working.
- **Traceroute**—The AFTR now generates and forwards traceroute packets over the DS-Lite tunnel.



NOTE: No additional CLI configuration is necessary to use the new functionality.

Monitoring Softwire Statistics

Purpose You can review softwire global statistics by using the **show services softwire** or **show services softwire statistics** command.

```
Action user@host# show services software
Interface: sp-0/0/0, Service set: sset
Software Direction Flow count
2001:0:0:1::1 -> 1001::1 I 3

user@host# show services software statistics
DS-Lite Statistics:
Service PIC Name: :sp-0/0/0
Statistics
-----
Softwires Created :2
Softwires Deleted :1
Softwires Flows Created :2
Softwires Flows Deleted :1
Slow Path Packets Processed :2
Fast Path Packets Processed :274240
Fast Path Packets Encapsulated :583337
Rule Match Failed :0
Rule Match Succeeded :2
IPv6 Packets Fragmented :0
Transient Errors
-----
Flow Creation Failed - Retry :0
Slow Path Failed - Retry :0
Errors
-----
Software Creation Failed :0
Flow Creation Failed :0
Slow Path Failed :0
Packet not IPv4-in-IPv6 :0
IPv6 Fragmentation Error :0
Slow Path Failed - IPv6 Next Header Offset :0
Decapsulated Packet not IPv4 :0
Fast Path Failed - IPv6 Next Header Offset :0
No Software ID :0
No Flow Extension :0
Flow Limit Exceeded :0
6rd Statistics:
Service PIC Name :sp-0/0/0
Statistics
-----
Softwires Created :0
Softwires Deleted :0
Softwires Flows Created :0
Softwires Flows Deleted :0
Slow Path Packets Processed :0
Fast Path Packets Processed :0
Fast Path Packets Encapsulated :0
Rule Match Failed :0
Rule Match Succeeded :0
Transient Errors
-----
Flow Creation Failed - Retry :0
Slow Path Failed - Retry :0
Errors
-----
Software Creation Failed :0
Flow Creation Failed :0
Slow Path Failed :0
Packet not IPv6-in-IPv4 :0
```

Slow Path Failed - IPv6 Next Header Offset :0
 Decapsulated Packet not IPv6 :0
 Encapsulation Failed - No packet memory :0
 No Softwire ID :0
 No Flow Extension :0
 ICMPv4 Dropped Packets :0

Monitoring CGN, Stateful Firewall, and Softwire Flows

Purpose Use the following commands to check the creation of the softwires, pre-NAT flows, and post-NAT flows. Output can be filtered using more specific fields such as AFTR or B4 address or both for DS-Lite, and softwire-concentrator or softwire-initiator or both for 6rd.

- [show services stateful-firewall flows](#)
- [show services softwire flows](#)

Action user@host# **show services stateful-firewall flows**
 Interface: sp-0/1/0, Service set: dslite-svc-set2

Flow				State	Dir	Frm count
TCP	200.200.200.2:80	->	44.44.44.1:1025	Forward	O	219942
	NAT dest		44.44.44.1:1025	->	20.20.1.4:1025	
	Softwire		2001::2	->	1001::1	
TCP	20.20.1.2:1025	->	200.200.200.2:80	Forward	I	110244
	NAT source		20.20.1.2:1025	->	44.44.44.1:1024	
	Softwire		2001::2	->	1001::1	
TCP	200.200.200.2:80	->	44.44.44.1:1024	Forward	O	219140
	NAT dest		44.44.44.1:1024	->	20.20.1.2:1025	
	Softwire		2001::2	->	1001::1	
DS-LITE	2001::2	->	1001::1	Forward	I	988729
TCP	200.200.200.2:80	->	44.44.44.1:1026	Forward	O	218906
	NAT dest		44.44.44.1:1026	->	20.20.1.3:1025	
	Softwire		2001::2	->	1001::1	
TCP	20.20.1.3:1025	->	200.200.200.2:80	Forward	I	110303
	NAT source		20.20.1.3:1025	->	44.44.44.1:1026	
	Softwire		2001::2	->	1001::1	
TCP	20.20.1.4:1025	->	200.200.200.2:80	Forward	I	110944
	NAT source		20.20.1.4:1025	->	44.44.44.1:1025	
	Softwire		2001::2	->	1001::1	

- Related Documentation**
- *NAT Objects MIB in SNMP MIBs and Traps Reference*
 - *Network Address Translation Resources—Monitoring MIB in SNMP MIBs and Traps Reference*
 - *Juniper Networks Enterprise-Specific NAT Traps on SRX Series Services Gateways in SNMP MIBs and Traps Reference*

PART 5

Enabling Traffic to Pass Securely Using ALGs

- [ALG Overview on page 277](#)
- [ALGs Configuration Overview on page 303](#)

CHAPTER 24

ALG Overview

- [ALG Descriptions on page 277](#)
- [ALGs Available by Default for Junos OS Address Aware NAT on page 300](#)

ALG Descriptions

This topic describes the Application Layer Gateways (ALGs) supported by Junos OS. ALG support includes managing pinholes and parent-child relationships for the supported ALGs. This topic includes the following sections:

- [Supported ALGs on page 277](#)
- [ALG Support Details on page 278](#)
- [Juniper Networks Defaults on page 288](#)
- [Examples: Referencing the Preset Statement from the Junos Default Group on page 298](#)

Supported ALGs

Table 13 on page 277 lists ALGs supported by Junos OS.

Table 13: ALGs Supported by Junos OS

ALGs Supported	v4 - v4	v4 - v6	v6 - v6	DS-Lite
Basic TCP ALG	Yes	Yes	Yes	Yes
Basic UPD ALG	Yes	Yes	Yes	Yes
BOOTP	Yes	No	No	No
DCE RPC Services	Yes	No	No	No
DNS	Yes	Yes	No	No
FTP	Yes	No	No	Yes
H323	Yes	No	No	No
ICMP	Yes	Yes	Yes	Yes

Table 13: ALGs Supported by Junos OS (*continued*)

ALGs Supported	v4 - v4	v4 - v6	v6 - v6	DS-Lite
IIOIP	Yes	No	No	No
IP	Yes	No	No	No
NETBIOS	Yes	No	No	No
NETSHOW	Yes	No	No	No
PPTP	Yes	No	No	Yes
REALAUDIO	Yes	No	No	No
Sun RPC and RPC Port Map Services	Yes	No	No	No
RTSP	Yes	No	No	Yes
SIP	Yes	No	No	No
SNMP	Yes	No	No	No
SQLNET	Yes	No	No	No
TFTP	Yes	No	No	Yes
Traceroute	Yes	Yes	No	Yes
Unix Remote Shell Service	Yes	No	No	No
WINFrame	Yes	No	No	No

ALG Support Details

This section includes details about the ALGs. It includes the following:

- [Basic TCP ALG on page 279](#)
- [Basic UDP ALG on page 279](#)
- [BOOTP on page 280](#)
- [DCE RPC Services on page 280](#)
- [DNS on page 280](#)
- [FTP on page 280](#)
- [H323 on page 281](#)
- [ICMP on page 282](#)
- [IIOIP on page 282](#)

- [IP on page 282](#)
- [NetBIOS on page 282](#)
- [NetShow on page 283](#)
- [ONC RPC Services on page 283](#)
- [PPTP on page 283](#)
- [RealAudio on page 283](#)
- [Sun RPC and RPC Portmap Services on page 284](#)
- [RTSP on page 285](#)
- [SIP on page 286](#)
- [SNMP on page 286](#)
- [SQLNet on page 286](#)
- [TFTP on page 287](#)
- [Traceroute on page 287](#)
- [UNIX Remote-Shell Services on page 287](#)
- [Winframe on page 288](#)

Basic TCP ALG

This ALG performs basic sanity checking on TCP packets. If it finds errors, it generates the following anomaly events and system log messages:

- TCP source or destination port zero
- TCP header length check failed
- TCP sequence number zero and no flags are set
- TCP sequence number zero and FIN/PSH/RST flags are set
- TCP FIN/RST or SYN(URG|FIN|RST) flags are set

The TCP ALG performs the following steps:

1. When the router receives a SYN packet, the ALG creates TCP forward and reverse flows and groups them in a *conversation*. It tracks the TCP three-way handshake.
2. The SYN-defense mechanism tracks the TCP connection establishment state. It expects the TCP session to be established within a small time interval (currently 4 seconds). If the TCP three-way handshake is not established in that period, the session is terminated.
3. A keepalive mechanism detects TCP sessions with nonresponsive endpoints.
4. ICMP errors are allowed only if there is a flow that matches the selector information specified in the ICMP data.

Basic UDP ALG

This ALG performs basic sanity checking on UDP headers. If it finds errors, it generates the following anomaly events and system log messages:

- UDP source or destination port 0
- UDP header length check failed

The UDP ALG performs the following steps:

1. When it receives the first packet, the ALG creates bidirectional flows to accept forward and reverse UDP session traffic.
2. If the session is idle for more than the maximum allowed idle time (the default is 30 seconds), the flows are deleted.
3. ICMP errors are allowed only if there is a flow that matches the selector information specified in the ICMP data.

BOOTP

The Bootstrap Protocol (BOOTP) client retrieves its networking information from a server across the network. It sends out a general broadcast message to request the information, which is returned by the BOOTP server. For the protocol specification, see <ftp://ftp.isi.edu/in-notes/rfc951.txt>.

Stateful firewall support requires that you configure the BOOTP ALG on UDP server port 67 and client port 68. If the client sends a broadcast message, you should configure the broadcast address in the **from** statement of the service rule. Network Address Translation (NAT) is not performed on the BOOTP traffic, even if the NAT rule matches the traffic. If the BOOTP relay feature is activated on the router, the remote BOOTP server is assumed to assign addresses for clients masked by NAT translation.

DCE RPC Services

Distributed Computing Environment (DCE) Remote Procedure Call (RPC) services are mainly used by Microsoft applications. The ALG uses well-known TCP port 135 for port mapping services, and uses the universal unique identifier (UUID) instead of the program number to identify protocols. The main application-based DCE RPC is the Microsoft Exchange Protocol.

Support for stateful firewall and NAT services requires that you configure the DCE RPC portmap ALG on TCP port 135. The DCE RPC ALG uses the TCP protocol with application-specific UUIDs.

DNS

The Domain Name Service (DNS) ALG handles data associated with locating and translating domain names into IP addresses. The ALG typically runs on port 53. The ALG monitors DNS query and reply packets and supports only UDP traffic. The ALG does not support payload translations. The DNS ALG will only close the session when a reply is received or an idle timeout is reached.

FTP

FTP is the File Transfer Protocol, specified in RFC 959. In addition to the main control connection, data connections are also made for any data transfer between the client

and the server; and the host, port, and direction are negotiated through the control channel.

For non-passive-mode FTP, Junos OS stateful firewall service scans the client-to-server application data for the PORT command, which provides the IP address and port number to which the server connects. For passive-mode FTP, Junos OS stateful firewall service scans the client-to-server application data for the PASV command and then scans the server-to-client responses for the 227 response, which contains the IP address and port number to which the client connects.

There is an additional complication: FTP represents these addresses and port numbers in ASCII. As a result, when addresses and ports are rewritten, the TCP sequence number might be changed, and thereafter the NAT service needs to maintain this delta in SEQ and ACK numbers by performing sequence NAT on all subsequent packets.

Support for stateful firewall and NAT services requires that you configure the FTP ALG on TCP port 21 to enable the FTP control protocol. The ALG performs the following tasks:

- Automatically allocates data ports and firewall permissions for dynamic data connection
- Creates flows for the dynamically negotiated data connection
- Monitors the control connection in both active and passive modes
- Rewrites the control packets with the appropriate NAT address and port information

On MS-MPCs and MS-MICs, for passive FTP to work properly without FTP application layer gateway (ALG) enabled (by not specifying the **application junos-ftp** statement at the **[edit services stateful-firewall rule rule-name term term-name from]** and the **[edit services nat rule rule-name term term-name from]** hierarchy levels), you must enable the address pooling paired (APP) functionality enabled (by including the **address-pooling** statement at the **[edit services nat rule rule-name term term-name then translated]** hierarchy level). Such a configuration causes the data and control FTP sessions to receive the same NAT address.

H323

H323 is a suite of ITU protocols for audio and video conferencing and collaboration applications. H323 consists of H.225 call signaling protocols and H.245 control protocol for media communication. During H.225 negotiation, the endpoints create a call by exchanging call signaling messages on the control channel and negotiate a new control channel for H.245. A new control connection is created for H.245 messages. Messages are exchanged on the H.245 control channel to open media channels.

Stateful firewall monitors the H.225 control channel to open the H.245 control channel. After the H.245 channel is created, stateful firewall also monitors this channel for media channel information and allows the media traffic through the firewall.

H323 ALG supports static destination, static and dynamic source NAT by rewriting the appropriate addresses and ports in the H.225 and H.245 messages.

ICMP

The Internet Control Message Protocol (ICMP) is defined in RFC 792. The Junos OS stateful firewall service allows ICMP messages to be filtered by specific type or specific type code value. ICMP error packets that lack a specifically configured type and code are matched against any existing flow in the opposite direction to check for the legitimacy of the error packet. ICMP error packets that pass the filter matching are subject to NAT translation.

The ICMP ALG always tracks ping traffic statefully using the ICMP sequence number. Each echo reply is forwarded only if there is an echo request with the corresponding sequence number. For any ping flow, only 20 echo requests can be forwarded without receiving an echo reply. When you configure dynamic NAT, the PING packet identifier is translated to allow additional hosts in the NAT pool to use the same identifier.

Support for stateful firewall and NAT services requires that you configure the ICMP ALG if the protocol is needed. You can configure the ICMP type and code for additional filtering.

IIOP

The Oracle Application Server NameServer Internet Inter-ORB Protocol (IIOP). This ALG is used in Common Object Request Broker Architecture (CORBA) based on distributed computing. Even though CORBA and IIOP are Object Management Group (OMG) standards, there is no fixed port assigned for IIOP. Each vendor implementing CORBA chooses a port. Java Virtual machine uses port 1975 by default, while ORBIX uses port 3075 as a default.

Stateful firewall and NAT require ALG IIOP be configured for TCP port 1975 for Java VM IIOP, and 3075 for CORBA applications ORBIX, a CORBA framework from Iona Technologies.

IP

The IP ALG is used to create uni-directional flows only. In case of TCP traffic, it does not check the 3-way handshake process. This ALG is useful in case of stateful firewall only service sets, where it allows traffic to flow uni-directionally only. When configuring in conjunction with **match-direction input-output** it allows the return traffic to flow through the stateful firewall as well. Typical scenarios are static NAT, destination NAT or scenarios where traffic is expected to traverse the stateful firewall in the presence of asymmetric routing. The Junos IP ALG is not intended for use with NAT, which will cause matching traffic to be discarded through the creation of a drop flow.

NetBIOS

A NetBIOS ALG translates NetBIOS IP addresses and port numbers when NAT is used.

NetBIOS supports the TCP and UDP transport protocols. Support for stateful firewall and NAT services requires that you configure the NetBIOS ALG on UDP port 138 and TCP port 139.

NetShow

The Microsoft protocol ms-streaming is used by NetShow, the Microsoft media server. This protocol supports several transport protocols: TCP, UDP, and HTTP. The client starts a TCP connection on port 1755 and sends the PORT command to the server. The server then starts UDP on that port to the client. Support for stateful firewall and NAT services requires that you configure the NetShow ALG on UDP port 1755.

ONC RPC Services

Open Networks Computing (ONC) RPC services function similarly to DCE RCP services. However, the ONC RPC ALG uses TCP/UDP port 111 for port mapping services, and uses the program number to identify protocols rather than the UUID.

Support for stateful firewall and NAT services requires that you configure the ONC RPC portmap ALG on TCP port 111. The ONC RPC ALG uses the TCP protocol with application-specific program numbers.

PPTP

The Point-to-Point Tunneling Protocol (PPTP) ALG is a TCP-based ALG. PPTP allows the Point-to-Point Protocol (PPP) to be tunneled through an IP network. PPTP defines a client-server architecture, a PPTP Network Server, and a PPTP Access Concentrator. The PPTP ALG requires a control connection and a data tunnel. The control connection uses TCP to establish and disconnect PPP sessions, and runs on port 1723. The data tunnel carries PPP traffic in generic routing encapsulated (GRE) packets that are carried over IP.

RealAudio

Real Networks PNA protocol RealVideo is not a separate service. It is part of the RealPlayer and most likely uses another channel for video. The RealPlayer versions G2, 7, and 8 use PNA and RTSP. For this version to work, the ALG must allow both PNA(7070) and RTSP(554). For the media, the server selects from a range of UDP ports(6970 through 7170), or TCP port 7071, or HTTP. The client can be configured to use a particular port. The RealPlayer versions 4.0 and 5.0 use control channel 7070 media UDP ports 6970 through 7170, or TCP port 7071, or HTTP. RealAudio player version 3.0 uses control channel 7070 media, UDP ports 6770-7170, or TCP port 7071.

Real products use the ports and ranges of ports shown in [Table 14 on page 283](#).

Table 14: RealAudio Product Port Usage

Real Product	Port Usage
4.0 and 5.0 Servers/Players	Control channel (bidirectional) on TCP port 7070. Data channel from server to player on TCP port 7070 or UDP port 6970-7170.
4.0 and 5.0 Servers/Encoders	Control channel (bidirectional) on TCP port 7070. Data channel from encoder or server on TCP port 7070.
G2 Servers/Players	Control channel (bidirectional) on TCP port 80, 554, 7070, or 8080. Data channel from server to player on TCP port 80, 554, 7070, 8080 or UDP port 6970-32,000.

Table 14: RealAudio Product Port Usage (*continued*)

Real Product	Port Usage
G2 Server/3.1, and 5.x Encoders	Control channel (bidirectional) on TCP port 7070. Data channel from encoder to server on TCP port 7070.
G2 Server/G2 Producer	Control channel (bidirectional) on TCP port 4040. Data channel from encoder to server on TCP port 4040 and UDP port 6970-32,000.
2 Server/G2 Producer (TCP ONLY)	Control channel (bidirectional) on TCP port 4040 Data channel from encoder to server on TCP port 4040. Note: TCP-ONLY option available in version 6.1 or above.



NOTE: RealAudio was the original protocol by RealPlayers. Newer versions of RealPlayer use RTSP. Stateful firewall and NAT require ALG RealAudio to be programmed on TCP port 7070.

Sun RPC and RPC Portmap Services

The Remote Procedure Call (RPC) ALG uses well-known ports TCP 111 and UDP 111 for port mapping, which dynamically assigns and opens ports for RPC services. The RPC Portmap ALG keeps track of port requests and dynamically opens the firewall for these requested ports. The RPC ALG can further restrict the RPC protocol by specifying allowed program numbers.

The ALG includes the RPC services listed in [Table 15 on page 284](#).

Table 15: Supported RPC Services

Name	Description	Comments
rpc-mountd	Network File Server (NFS) mount daemon; for details, see the UNIX man page for rpc.mountd(8) .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050).
rpc-nfsprog	Used as part of NFS. For details, see RFC 1094. See also RFC1813 for NFS v3.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050).
rpc-nisplus	Network Information Service Plus (NIS+), designed to replace NIS; it is a default naming service for Sun Solaris and is not related to the old NIS. No protocol information is available.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050).
rpc-nlockmgr	Network lock manager.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, rpc-nlockmgr service can be allowed or blocked based on RPC program 100021.
rpc-pcnfsd	Kernel statistics server. For details, see the UNIX man pages for rstatd and rpc.rstatd .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, rpc-rstat service can be allowed or blocked based on RPC program 150001.

Table 15: Supported RPC Services (*continued*)

Name	Description	Comments
rpc-rwall	Used to write a message to users; for details, see the UNIX man page for rpc.rwalld .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, rpc-rwall service can be allowed or blocked based on RPC program 150008.
rpc-ybind	NIS binding process. For details, see the UNIX man page for ybind .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, rpc-ybind service can be allowed or blocked based on RPC program 100007.
rpc-yppasswd	NIS password server. For details, see the UNIX man page for yppasswd .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, rpc-yppasswd service can be allowed or blocked based on RPC program 100009.
rpc-ybserv	NIS server. For details, see the UNIX man page for ybserv .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, rpc-ybserv service can be allowed or blocked based on RPC program 100004.
rpc-ypupdated	Network updating tool.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, rpc-ypupdated service can be allowed or blocked based on RPC program 100028.
rpc-ypxfrd	NIS map transfer server. For details, see the UNIX man page for rpc.ypxfrd .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, rpc-ypxfrd service can be allowed or blocked based on RPC program 100069.

Support for stateful firewall and NAT services that use port mapping requires that you configure the RPC portmap ALG on TCP/UDP destination port 111 and the RPC ALG for both TCP and UDP. You can specify one or more **rpc-program-number** values to further restrict allowed RPC protocols.

RTSP

The Real-Time Streaming Protocol (RTSP) controls the delivery of data with real-time properties such as audio and video. The streams controlled by RTSP can use RTP, but it is not required. Media can be transmitted on the same RTSP control stream. This is an HTTP-like text-based protocol, but client and server maintain session information. A session is established using the SETUP message and terminated using the TEARDOWN message. The transport (the media protocol, address, and port numbers) is negotiated in the setup and the setup-response.

Support for stateful firewall and NAT services requires that you configure the RTSP ALG for TCP port 554.

The ALG monitors the control connection, opens flows dynamically for media (RTP/RTSP) streams, and performs NAT address and port rewrites.

SIP

The Session Initiation Protocol (SIP) is an application layer protocol that can establish, maintain, and terminate media sessions. It is a widely used voice over IP (VoIP) signaling protocol. The SIP ALG monitors SIP traffic and dynamically creates and manages pinholes on the signaling and media paths. The ALG only allows packets with the correct permissions. The SIP ALG also performs the following functions:

- Manages parent-child session relationships.
- Enforces security policies.
- Manages pinholes for VoIP traffic.

The SIP ALG supports the following features:

- Stateful firewall
- Static source NAT
- Dynamic address only source NAT
- Network Address Port Translation (NAPT)



NOTE: SIP sessions are limited to 12 hours (720 minutes) for NAT processing on the MS-MIC and MS-MPC interface cards. There is no time limit for SIP sessions on the MS-DPC.

SNMP

SNMP is a communication protocol for managing TCP/IP networks, including both individual network devices and aggregated devices. The protocol is defined by RFC 1157. SNMP runs on top of UDP.

The Junos OS stateful firewall service implements the SNMP ALG to inspect the SNMP type. SNMP does not enforce stateful flow. Each SNMP type needs to be specifically enabled. Full SNMP support of stateful firewall services requires that you configure the SNMP ALG on UDP port 161. This enables the SNMP **get** and **get-next** commands, as well as their response traffic in the reverse direction: UDP port 161 enables the SNMP **get-response** command. If SNMP traps are permitted, you can configure them on UDP port 162, enabling the SNMP **trap** command.

SQLNet

The SQLNet protocol is used by Oracle SQL servers to execute SQL commands from clients, including load balancing and application-specific services.

Support of stateful firewall and NAT services requires that you configure the SQLNet ALG for TCP port 1521.

The ALG monitors the control packets, opens flows dynamically for data traffic, and performs NAT address and port rewrites.

TFTP

The Trivial File Transfer Protocol (TFTP) is specified in RFC 1350. The initial TFTP requests are sent to UDP destination port 69. Additional flows can be created to **get** or **put** individual files. Support of stateful firewall and NAT services requires that you configure the TFTP ALG for UDP destination port 69.

Traceroute

Traceroute is a tool for displaying the route that packets take to a network host. It uses the IP time-to-live (TTL) field to trigger ICMP time-exceeded messages from routers or gateways. It sends UDP datagrams to destination ports that are believed to be not in use; destination ports are numbered using the formula: $+ n\text{hops} - 1$. The default base port is 33434. To support traceroute through the firewall, two types of traffic must be passed through:

1. UDP probe packets (UDP destination port > 33000 , IP TTL < 30)
2. ICMP response packets (ICMP type time-exceeded)

When NAT is applied, the IP address and port within the ICMP error packet also must be changed.

Support of stateful firewall and NAT services requires you to configure the Traceroute ALG for UDP destination port 33434 to 33450. In addition, you can configure the TTL threshold to prevent UDP flood attacks with large TTL values.

UNIX Remote-Shell Services

Three protocols form the basis for UNIX remote-shell services:

- **Exec**—Remote command execution; enables a user on the client system to execute a command on the remote system. The first command from client (**rcmd**) to server (**rshd**) uses well-known TCP port 512. A second TCP connection can be opened at the request of **rcmd**. The client port number for the second connection is sent to the server as an ASCII string.
- **Login**—Better known as **rlogin**; uses well-known TCP port 513. For details, see RFC 1282. No special firewall processing is required.
- **Shell**—Remote command execution; enables a user on the client system to execute a command on the remote system. The first command from client (**rcmd**) to server (**rshd**) uses well-known TCP port 514. A second TCP connection can be opened at the request of **rcmd**. The client port number for the second connection is sent to the server as an ASCII string.

Support of stateful firewall services requires that you configure the Exec ALG on TCP port 512, the Login ALG on TCP port 513, and the Shell ALG on TCP port 514. NAT remote-shell services require that any dynamic source port assigned be within the port range 512 to 1023. If you configure a NAT pool, this port range is reserved exclusively for remote shell applications.

Winframe

WinFrame application server software provides access to virtually any Windows application, across any type of network connection to any type of client.

This protocol is mainly used by Citrix Windows applications.

Stateful firewall and NAT require the ALG Winframe to be configured on TCP destination port 1494 and UDP port 1604.

Juniper Networks Defaults

The Junos OS provides a default, hidden configuration group called **junos-defaults** that is automatically applied to the configuration of your router. The **junos-defaults** group contains preconfigured statements that contain predefined values for common applications. Some of the statements must be referenced to take effect, such as applications like FTP or Telnet. Other statements are applied automatically, such as terminal settings. All of the preconfigured statements begin with the reserved name **junos-**.



NOTE: You can override the Junos default configuration values, but you cannot delete or edit them. If you delete a configuration, the defaults return when a new configuration is added.

You cannot use the **apply-groups** statement with the Junos defaults group.

To view the full set of available preset statements from the Junos default group, issue the **show groups junos-defaults** configuration mode command. The following example displays the list of Junos default groups that use application protocols (ALGs).

```
user@host# show groups junos-defaults
applications {
  #
  # File Transfer Protocol
  #
  application junos-ftp {
    application-protocol ftp;
    protocol tcp;
    destination-port 21;
  }
  #
  # Trivial File Transfer Protocol
  #
  application junos-tftp {
    application-protocol tftp;
    protocol udp;
    destination-port 69;
  }
  #
  # RPC portmapper on TCP
  #
  application junos-rpc-portmap-tcp {
```

```
    application-protocol rpc-portmap;
    protocol tcp;
    destination-port 111;
}
#
# RPC portmapper on UDP
#
application junos-rpc-portmap-udp {
    application-protocol rpc-portmap;
    protocol udp;
    destination-port 111;
}
#
# SNMP get
#
application junos-snmp-get {
    application-protocol snmp;
    protocol udp;
    destination-port 161;
    snmp-command get;
}
#
# SNMP get next
#
application junos-snmp-get-next {
    application-protocol snmp;
    protocol udp;
    destination-port 161;
    snmp-command get-next;
}
#
# SNMP response
#
application junos-snmp-response {
    application-protocol snmp;
    protocol udp;
    source-port 161;
    snmp-command get-response;
}
#
# SNMP trap
#
application junos-snmp-trap {
    application-protocol snmp;
    protocol udp;
    destination-port 162;
    snmp-command trap;
}
#
# remote exec
#
application junos-rexec {
    application-protocol exec;
    protocol tcp;
    destination-port 512;
}
```

```
#
# remote login
#
application junos-rlogin {
    application-protocol shell;
    protocol tcp;
    destination-port 513;
}
#
# remote shell
#
application junos-rsh {
    application-protocol shell;
    protocol tcp;
    destination-port 514;
}
#
# Real Time Streaming Protocol
#
application junos-rtsp {
    application-protocol rtsp;
    protocol tcp;
    destination-port 554;
}
#
# Citrix windows application server protocol
# windows applications remotely on windows/non-windows clients
#
# citrix needs udp 1604 to be open
#
application junos-citrix-winframe {
    application-protocol winframe;
    protocol tcp;
    destination-port 1494;
}
application junos-citrix-winframe-udp {
    protocol udp;
    destination-port 1604;
}
#
# Oracle SQL servers use this protocol to execute sql commands
# from clients, load balance, use application-specific servers, etc
#
application junos-sqlnet {
    application-protocol sqlnet;
    protocol tcp;
    destination-port 1521;
}
#
# H.323 Protocol for audio/video conferencing
#
application junos-h323 {
    application-protocol h323;
    protocol tcp;
    destination-port 1720;
}
```

```
#
# Internet Inter-ORB Protocol - used for CORBA applications
# The ORB protocol in Java virtual machines uses port 1975 as default
#
application junos-iiop-java {
    application-protocol iiop;
    protocol tcp;
    destination-port 1975;
}
#
# Internet Inter-ORB Protocol - used for CORBA applications
# ORBIX is a CORBA framework from Iona Technologies that uses port
# 3075 as default
#
application junos-iiop-orbix {
    application-protocol iiop;
    protocol tcp;
    destination-port 3075;
}
#
# Real players use this protocol for real time streaming
# This was the original protocol for real players.
# RTSP is more widely used by real players
# but they still support realaudio.
#
application junos-realaudio {
    application-protocol realaudio;
    protocol tcp;
    destination-port 7070;
}
#
# traceroute application.
#
application junos-traceroute {
    application-protocol traceroute;
    protocol udp;
    destination-port 33435-33450;
    ttl-threshold 30;
}
#
# The full range of known RPC programs using UDP
# The program numbers can be more specific to certain applications.
#
application junos-rpc-services-udp {
    application-protocol rpc;
    protocol udp;
    rpc-program-number 100000-400000;
}
#
# The full range of known RPC programs using TCP
# The program numbers can be more specific to certain applications.
#
application junos-rpc-services-tcp {
    application-protocol rpc;
    protocol tcp;
    rpc-program-number 100000-400000;
```

```
}
#
# All ICMP traffic
# This can be made to be more restrictive by specifying ICMP type
# and code.
#
application junos-icmp-all {
    application-protocol icmp;
}
#
# Protocol used by Windows media server and windows media player
#
application junos-netshow {
    application-protocol netshow;
    protocol tcp;
    destination-port 1755;
}
#
# NetBIOS - networking protocol used on
# Windows networks name service port, both UDP and TCP
#
application junos-netbios-name-udp {
    application-protocol netbios;
    protocol udp;
    destination-port 137;
}
application junos-netbios-name-tcp {
    protocol tcp;
    destination-port 137;
}
#
# NetBIOS - networking protocol used on
# Windows networks datagram service port
#
application junos-netbios-datagram {
    application-protocol netbios;
    protocol udp;
    destination-port 138;
}
#
# NetBIOS - networking protocol used on
# Windows networks session service port
#
application junos-netbios-session {
    protocol tcp;
    destination-port 139;
}
#
# DCE-RPC portmapper on TCP
#
application junos-dce-rpc-portmap {
    application-protocol dce-rpc-portmap;
    protocol tcp;
    destination-port 135;
}
#
```

```
# DCE-RPC application on TCP sample UUID
# This application requires user to specify the UUID value
#
# application junos-dcerpc {
#   application-protocol dce-rpc;
#   protocol tcp;
#   #
#   # UUID also needs to be defined as shown below
#   UUID 11223344 22334455 33445566 44556677;
#   #
# }
#
# ms-exchange needs these 3 UUIDs
#
application junos-dcerpc-endpoint-mapper-service {
  application-protocol dce-rpc;
  protocol tcp;
  uuid e1af8308-5d1f-11c9-91a4-08002b14a0fa;
}
application junos-dcerpc-msexchange-directory-rfr {
  application-protocol dce-rpc;
  protocol tcp;
  uuid 1544f5e0-613c-11d1-93df-00c04fd7bd09;
}
application junos-dcerpc-msexchange-information-store {
  application-protocol dce-rpc;
  protocol tcp;
  uuid a4f1db00-ca47-1067-b31f-00dd010662da;
}
application junos-ssh {
  protocol tcp;
  destination-port 22;
}
application junos-telnet {
  protocol tcp;
  destination-port 23;
}
application junos-smtp {
  protocol tcp;
  destination-port 25;
}
application junos-dns-udp {
  protocol udp;
  destination-port 53;
}
application junos-dns-tcp {
  protocol tcp;
  destination-port 53;
}
application junos-tacacs {
  protocol tcp;
  destination-port 49;
}
# TACACS Database Service
application junos-tacacs-ds {
  protocol tcp;
```

```
        destination-port 65;
    }
    application junos-dhcp-client {
        protocol udp;
        destination-port 68;
    }
    application junos-dhcp-server {
        protocol udp;
        destination-port 67;
    }
    application junos-bootpc {
        protocol udp;
        destination-port 68;
    }
    application junos-bootps {
        protocol udp;
        destination-port 67;
    }
    application junos-finger {
        protocol tcp;
        destination-port 79;
    }
    application junos-http {
        protocol tcp;
        destination-port 80;
    }
    application junos-https {
        protocol tcp;
        destination-port 443;
    }
    application junos-pop3 {
        protocol tcp;
        destination-port 110;
    }
    application junos-ident {
        protocol tcp;
        destination-port 113;
    }
    application junos-nntp {
        protocol tcp;
        destination-port 119;
    }
    application junos-ntp {
        protocol udp;
        destination-port 123;
    }
    application junos-imap {
        protocol tcp;
        destination-port 143;
    }
    application junos-imaps {
        protocol tcp;
        destination-port 993;
    }
    application junos-bgp {
        protocol tcp;
```



```
        destination-port 179;
    }
    application junos-ldap {
        protocol tcp;
        destination-port 389;
    }
    application junos-snpp {
        protocol tcp;
        destination-port 444;
    }
    application junos-biff {
        protocol udp;
        destination-port 512;
    }
    # UNIX who
    application junos-who {
        protocol udp;
        destination-port 513;
    }
    application junos-syslog {
        protocol udp;
        destination-port 514;
    }
    # line printer daemon, printer, spooler
    application junos-printer {
        protocol tcp;
        destination-port 515;
    }
    # UNIX talk
    application junos-talk-tcp {
        protocol tcp;
        destination-port 517;
    }
    application junos-talk-udp {
        protocol udp;
        destination-port 517;
    }
    application junos-ntalk {
        protocol udp;
        destination-port 518;
    }
    application junos-rip {
        protocol udp;
        destination-port 520;
    }
    # INA sanctioned RADIUS port numbers
    application junos-radius {
        protocol udp;
        destination-port 1812;
    }
    application junos-radacct {
        protocol udp;
        destination-port 1813;
    }
    application junos-nfsd-tcp {
        protocol tcp;
```

```
        destination-port 2049;
    }
    application junos-nfsd-udp {
        protocol udp;
        destination-port 2049;
    }
    application junos-cvsspserver {
        protocol tcp;
        destination-port 2401;
    }
    #
    # Label Distribution Protocol
    #
    application junos-ldp-tcp {
        protocol tcp;
        destination-port 646;
    }
    application junos-ldp-udp {
        protocol udp;
        destination-port 646;
    }
    #
    # JUNOScript and JUNOScope management
    #
    application junos-xnm-ssl {
        protocol tcp;
        destination-port 3220;
    }
    application junos-xnm-clear-text {
        protocol tcp;
        destination-port 3221;
    }
    #
    # IPsec tunnel
    #
    application junos-ipsec-esp {
        protocol esp;
    }
    application junos-ike {
        protocol udp;
        destination-port 500;
    }
    #
    # 'junos-algs-outbound' defines a set of all applications
    # requiring an ALG. Useful for defining rule to the the public
    # internet allowing private network users to use all JUNOS OS
    # supported ALGs initiated from the private network.
    #
    # NOTE: the contents of this set might grow in future JUNOS OS versions.
    #
    application-set junos-algs-outbound {
        application junos-ftp;
        application junos-tftp;
        application junos-rpc-portmap-tcp;
        application junos-rpc-portmap-udp;
        application junos-snmp-get;
```

```

application junos-snmp-get-next;
application junos-snmp-response;
application junos-snmp-trap;
application junos-rexec;
application junos-rlogin;
application junos-rsh;
application junos-rtsp;
application junos-citrix-winfile;
application junos-citrix-winfile-udp;
application junos-sqlnet;
application junos-h323;
application junos-iiop-java;
application junos-iiop-orbix;
application junos-realaudio;
application junos-traceroute;
application junos-rpc-services-udp;
application junos-rpc-services-tcp;
application junos-icmp-all;
application junos-netshow;
application junos-netbios-name-udp;
application junos-netbios-datagram;
application junos-dcerpc-endpoint-mapper-service;
application junos-dcerpc-msexchange-directory-rfr;
application junos-dcerpc-msexchange-information-store;
}
#
# 'junos-management-inbound' represents the group of applications
# that might need access the router from public network for
# for management purposes.
#
# Set is intended for a UI to display management choices.
#
# NOTE: It is not recommended the user to use the entire set
# directly in a firewall rule and open up firewall to all
# of these applications. Also, the user should always
# specify the source and destination prefixes when using
# each application.
#
# NOTE: the contents of this set may grow in future JUNOS versions.
#
application-set junos-management-inbound {
    application junos-snmp-get;
    application junos-snmp-get-next;
    application junos-snmp-response;
    application junos-snmp-trap;
    application junos-ssh;
    application junos-telnet;
    application junos-http;
    application junos-https;
    application junos-xnm-ssl;
    application junos-xnm-clear-text;
}
#
# 'junos-routing-inbound' represents routing protocols that might
# need to access the router from public network.
#

```

```
# Set is intended for a UI to display routing involvement choices.
#
# NOTE: It is not recommended the user to use the entire set
# directly in a firewall rule and open up firewall to all
# of these applications. Also, the user should always
# specify the source and destination prefixes when using
# each application.
#
# NOTE: the contents of this set might grow in future JUNOS OS versions.
#
application-set junos-routing-inbound {
    application junos-bgp;
    application junos-rip;
    application junos-ldp-tcp;
    application junos-ldp-udp;
}
```

To reference statements available from the **junos-defaults** group, include the selected **junos-default-name** statement at the applicable hierarchy level. To configure application protocols, see [“Configuring Application Protocol Properties” on page 303](#); for details about a specific protocol, see [“ALG Descriptions” on page 277](#).

Examples: Referencing the Preset Statement from the Junos Default Group

The following example is a preset statement from the Junos default groups that is available for FTP in a stateful firewall:

```
[edit]
groups {
  junos-defaults {
    applications {
      application junos-ftp { # Use FTP default configuration
        application-protocol ftp;
        protocol tcp;
        destination-port 21;
      }
    }
  }
}
```

To reference a preset Junos default statement from the Junos default groups, include the **junos-default-name** statement at the applicable hierarchy level. For example, to reference the Junos default statement for FTP in a stateful firewall, include the **junos-ftp** statement at the **[edit services stateful-firewall rule rule-name term term-name from applications]** hierarchy level.

```
[edit]
services {
  stateful-firewall {
    rule my-rule {
      term my-term {
        from {
          applications junos-ftp; #Reference predefined statement, junos-ftp,
        }
      }
    }
  }
}
```

```

    }
  }

```

The following example shows configuration of the default Junos IP ALG:

```

[edit]
services {
  stateful-firewall {
    rule r1 {
      match-direction input;
      term t1 {
        from {
          applications junos-ip;
        }
        then {
          accept;
          syslog;
        }
      }
    }
  }
}

```

If you configure the IP ALG in the stateful firewall rule, it is matched by any IP traffic, but if there is any other more specific application that matches the same traffic, the IP ALG will not be matched. For example, in the following configuration, both the ICMP ALG and the IP ALG are configured, but traffic is matched for ICMP packets, because it is the more specific match.

```

[edit]
services {
  stateful-firewall {
    rule r1 {
      match-direction input;
      term t1 {
        from {
          applications [ junos-ip junos-icmp-all ];
        }
        then {
          accept;
          syslog;
        }
      }
    }
  }
}

```

Related Documentation

- *OBSOLETE Junos Default Groups*
- [Configuring Application Sets on page 303](#)
- [Configuring Application Protocol Properties on page 303](#)

ALGs Available by Default for Junos OS Address Aware NAT

The following application-level gateways (ALGs) listed in [Table 12 on page 141](#) are supported for NAT processing on the listed platforms.

To view the implementation details (port, protocol, and so on) for these Junos OS default applications, locate the Junos OS Default ALG Name in the table and then look up the listed name in the **groups**. For example, for details about TFTP, look up **junos-tftp** as shown.



TIP: The Junos OS provides the **junos-alg**, which enables other ALGs to function by handling ALG registrations, causing slow path packets to flow through registered ALGs, and transferring ALG events to the ALG plug-ins. The **junos-alg** ALG is automatically available on the MS-MPC and MS-MIC platforms and does not require further configuration.

```
user@host# show groups junos-defaults applications application junos-tftp
application-protocol tftp;
protocol udp;
destination-port 69;
```

Table 16: ALGs Available by Default

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
Basic TCP ALG	yes	yes	NOTE: Specific Junos ALGs are not supported. However, a feature called TCP tracker, available by default, performs segment ordering and retransmit and connection tracking, validations for TCP connections.
Basic UDP ALG	yes	yes	NOTE: TCP tracker performs limited integrity and validation checks for UDP.
BOOTP	yes	no	<ul style="list-style-type: none"> junos-bootpc junos-bootps
DCE RPC Services	yes	yes	<ul style="list-style-type: none"> junos-dce-rpc-portmap junos-dcerpc-endpoint-mapper-service junos-dcerpc-msexchange-directory-nsp junos-dcerpc-msexchange-directory-rfr junos-dcerpc-msexchange-information-store
DNS	yes	yes	<ul style="list-style-type: none"> junos-dns-tcp junos-dns-udp
FTP	yes	yes	<ul style="list-style-type: none"> junos-ftp
H323	yes	no	<ul style="list-style-type: none"> junos-h323

Table 16: ALGs Available by Default (*continued*)

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
ICMP	yes	yes <i>NOTE:</i> ICMP messages are handled by default, but PING ALG support is not provided.	<ul style="list-style-type: none"> • <code>junos-icmp-all</code> • <code>junos-icmp-ping</code>
IIOIP	yes	no	<ul style="list-style-type: none"> • <code>junos-iioip-java</code> • <code>junos-iioip-orbix</code>
IP	yes	The TCP tracker, available by default on these platforms, performs limited integrity and validation checks.	<ul style="list-style-type: none"> • <code>junos-ip</code>
NETBIOS	yes	no	<ul style="list-style-type: none"> • <code>junos-netbios-datagram</code> • <code>junos-netbios-name-tcp</code> • <code>junos-netbios-name-udp</code> • <code>junos-netbios-session</code>
NETSHOW	yes	no	<ul style="list-style-type: none"> • <code>junos-netshow</code>
PPTP	yes	yes	<ul style="list-style-type: none"> • <code>junos-pptp</code>
REALAUDIO	yes	no	<ul style="list-style-type: none"> • <code>junos-realaudio</code>
Sun RPC and RPC Port Map Services	yes	yes	<ul style="list-style-type: none"> • <code>junos-rpc-portmap-tcp</code> • <code>junos-rpc-portmap-udp</code>
RTSP	yes	yes	<ul style="list-style-type: none"> • <code>junos-rtsp</code>
SIP	yes	Yes	<ul style="list-style-type: none"> • <code>junos-sip</code> <p>The SIP <code>callid</code> is <i>not</i> translated in <code>register</code> messages.</p> <p><i>NOTE:</i> SIP sessions are limited to 12 hours (720 minutes) for NAT processing on the MS-MIC and MS-MPC interface cards. There is no time limit for SIP sessions on the MS-DPC.</p>
SNMP	yes	No	<ul style="list-style-type: none"> • <code>junos-snmp-get</code> • <code>junos-snmp-get-next</code> • <code>junos-snmp-response</code> <code>junos-snmp-trap</code>
SQLNET	yes	yes	<ul style="list-style-type: none"> • <code>junos-sqlnet</code>
TFTP	yes	yes	<ul style="list-style-type: none"> • <code>junos-tftp</code>

Table 16: ALGs Available by Default (*continued*)

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
Traceroute	yes	no	<ul style="list-style-type: none"> • <code>junos-traceroute</code>
Unix Remote Shell Service	yes	Yes	<ul style="list-style-type: none"> • <code>junos-rsh</code>
WINFrame	yes	No	<ul style="list-style-type: none"> • <code>junos-citrix-winframe</code> • <code>junos-citrix-winframe-udp</code>
TALK-UDP	No	Yes	<ul style="list-style-type: none"> • <code>junos-talk-udp</code>
MS RPC	No	Yes	<ul style="list-style-type: none"> • <code>junos-rpc-portmap-tcp</code> • <code>junos-rpc-portmap-udp</code> • <code>junos-rpc-services-tcp</code> • <code>junos-rpc-services-udp</code>

Related Documentation • [ALG Descriptions on page 277](#)

ALGs Configuration Overview

- [Configuring Application Sets on page 303](#)
- [Configuring Application Protocol Properties on page 303](#)
- [Examples: Configuring Application Protocols on page 321](#)
- [ICMP, Ping, and Traceroute ALGs for MS-MICs and MS-MPCs on page 322](#)
- [Monitoring Port Control Protocol Operations on page 323](#)

Configuring Application Sets

You can group the applications you have defined into a named object by including the **application-set** statement at the **[edit applications]** hierarchy level with an **application** statement for each application:

```
[edit applications]
  application-set application-set-name {
    application application;
  }
```

For an example of a typical application set, see “[Examples: Configuring Application Protocols](#)” on page 321.

Related Documentation

- [ALG Descriptions on page 277](#)
- [Configuring Application Protocol Properties on page 303](#)
- [Examples: Configuring Application Protocols on page 321](#)
- [Verifying the Output of ALG Sessions](#)

Configuring Application Protocol Properties

To configure application properties, include the **application** statement at the **[edit applications]** hierarchy level:

```
[edit applications]
  application application-name {
    application-protocol protocol-name;
    destination-port port-number;
    icmp-code value;
    icmp-type value;
```

```

inactivity-timeout value;
protocol type;
rpc-program-number number;
snmp-command command;
source-port port-number;
ttl-threshold value;
uuid hex-value;
}

```

You can group application objects by configuring the **application-set** statement; for more information, see “Configuring Application Sets” on page 303.

This section includes the following tasks for configuring applications:

- [Configuring an Application Protocol on page 304](#)
- [Configuring the Network Protocol on page 306](#)
- [Configuring the ICMP Code and Type on page 307](#)
- [Configuring Source and Destination Ports on page 309](#)
- [Configuring the Inactivity Timeout Period on page 312](#)
- [Configuring SIP on page 312](#)
- [Configuring an SNMP Command for Packet Matching on page 320](#)
- [Configuring an RPC Program Number on page 320](#)
- [Configuring the TTL Threshold on page 320](#)
- [Configuring a Universal Unique Identifier on page 320](#)

Configuring an Application Protocol

The **application-protocol** statement allows you to specify which of the supported application protocols (ALGs) to configure and include in an application set for service processing. To configure application protocols, include the **application-protocol** statement at the **[edit applications application application-name]** hierarchy level:

```

[edit applications application application-name]
application-protocol protocol-name;

```

[Table 17 on page 304](#) shows the list of supported protocols. For more information about specific protocols, see “ALG Descriptions” on page 277.

Table 17: Application Protocols Supported by Services Interfaces

Protocol Name	CLI Value	Comments
Bootstrap protocol (BOOTP)	bootp	Supports BOOTP and dynamic host configuration protocol (DHCP).
Distributed Computing Environment (DCE) remote procedure call (RPC)	dce-rpc	Requires the protocol statement to have the value udp or tcp . Requires a uuid value. You cannot specify destination-port or source-port values.
DCE RPC portmap	dce-rpc-portmap	Requires the protocol statement to have the value udp or tcp . Requires a destination-port value.

Table 17: Application Protocols Supported by Services Interfaces (*continued*)

Protocol Name	CLI Value	Comments
Domain Name System (DNS)	dns	Requires the protocol statement to have the value udp . This application protocol closes the DNS flow as soon as the DNS response is received.
Exec	exec	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.
FTP	ftp	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.
H.323	h323	—
Internet Control Message Protocol (ICMP)	icmp	Requires the protocol statement to have the value icmp or to be unspecified.
Internet Inter-ORB Protocol	iiop	—
IP	ip	—
Login	login	—
NetBIOS	netbios	Requires the protocol statement to have the value udp or to be unspecified. Requires a destination-port value.
NetShow	netshow	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.
Point-to-Point Tunneling Protocol	pptp	—
RealAudio	realaudio	—
Real-Time Streaming Protocol (RTSP)	rtsp	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.
RPC User Datagram Protocol (UDP) or TCP	rpc	Requires the protocol statement to have the value udp or tcp . Requires a rpc-program-number value. You cannot specify destination-port or source-port values.
RPC port mapping	rpc-portmap	Requires the protocol statement to have the value udp or tcp . Requires a destination-port value.
Shell	shell	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.
Session Initiation Protocol	sip	—
SNMP	snmp	Requires the protocol statement to have the value udp or to be unspecified. Requires a destination-port value.

Table 17: Application Protocols Supported by Services Interfaces (*continued*)

Protocol Name	CLI Value	Comments
SQLNet	sqlnet	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port or source-port value.
Talk Program	talk	
Trace route	traceroute	Requires the protocol statement to have the value udp or to be unspecified. Requires a destination-port value.
Trivial FTP (TFTP)	tftp	Requires the protocol statement to have the value udp or to be unspecified. Requires a destination-port value.
WinFrame	winframe	—



NOTE: You can configure application-level gateways (ALGs) for ICMP and trace route under stateful firewall, NAT, or CoS rules when twice NAT is configured in the same service set. These ALGs cannot be applied to flows created by the Packet Gateway Controller Protocol (PGCP). Twice NAT does not support any other ALGs. NAT applies only the IP address and TCP or UDP headers, but not the payload.

For more information about configuring twice NAT, see *Junos Address Aware Carrier Grade NAT and IPv6 Feature Guide*.

Related Documentation

- [ALGs Available by Default for Junos OS Address Aware NAT on page 141](#)

Configuring the Network Protocol

The **protocol** statement allows you to specify which of the supported network protocols to match in an application definition. To configure network protocols, include the **protocol** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
  protocol type;
```

You specify the protocol type as a numeric value; for the more commonly used protocols, text names are also supported in the command-line interface (CLI). [Table 18 on page 306](#) shows the list of the supported protocols.

Table 18: Network Protocols Supported by Services Interfaces

Network Protocol Type	CLI Value	Comments
IP Security (IPsec) authentication header (AH)	ah	—

Table 18: Network Protocols Supported by Services Interfaces (*continued*)

Network Protocol Type	CLI Value	Comments
External Gateway Protocol (EGP)	egp	—
IPsec Encapsulating Security Payload (ESP)	esp	—
Generic routing encapsulation (GR)	gre	—
ICMP	icmp	Requires an application-protocol value of icmp .
ICMPv6	icmp6	Requires an application-protocol value of icmp .
Internet Group Management Protocol (IGMP)	igmp	—
IP in IP	ipip	—
OSPF	ospf	—
Protocol Independent Multicast (PIM)	pim	—
Resource Reservation Protocol (RSVP)	rsvp	—
TCP	tcp	Requires a destination-port or source-port value unless you specify application-protocol rcp or dce-rcp .
UDP	udp	Requires a destination-port or source-port value unless you specify application-protocol rcp or dce-rcp .

For a complete list of possible numeric values, see RFC 1700, *Assigned Numbers (for the Internet Protocol Suite)*.



NOTE: IP version 6 (IPv6) is not supported as a network protocol in application definitions.

By default, the twice NAT feature can affect IP, TCP, and UDP headers embedded in the payload of ICMP error messages. You can include the **protocol tcp** and **protocol udp** statements with the application statement for twice NAT configurations. For more information about configuring twice NAT, see *Junos Address Aware Carrier Grade NAT and IPv6 Feature Guide*.

Configuring the ICMP Code and Type

The ICMP code and type provide additional specification, in conjunction with the network protocol, for packet matching in an application definition. To configure ICMP settings,

include the **icmp-code** and **icmp-type** statements at the **[edit applications application application-name]** hierarchy level:

```
[edit applications application application-name]
icmp-code value;
icmp-type value;
```

You can include only one ICMP code and type value. The **application-protocol** statement must have the value **icmp**. [Table 19 on page 308](#) shows the list of supported ICMP values.

Table 19: ICMP Codes and Types Supported by Services Interfaces

CLI Statement	Description
icmp-code	<p>This value or keyword provides more specific information than icmp-type. Because the value's meaning depends upon the associated icmp-type value, you must specify icmp-type along with icmp-code. For more information, see the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <p>parameter-problem: ip-header-bad (0), required-option-missing (1)</p> <p>redirect: redirect-for-host (1), redirect-for-network (0), redirect-for-tos-and-host (3), redirect-for-tos-and-net (2)</p> <p>time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0)</p> <p>unreachable: communication-prohibited-by-filtering (13), destination-host-prohibited (10), destination-host-unknown (7), destination-network-prohibited (9), destination-network-unknown (6), fragmentation-needed (4), host-precedence-violation (14), host-unreachable (1), host-unreachable-for-TOS (12), network-unreachable (0), network-unreachable-for-TOS (11), port-unreachable (3), precedence-cutoff-in-effect (15), protocol-unreachable (2), source-host-isolated (8), source-route-failed (5)</p>
icmp-type	<p>Normally, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port. For more information, see the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): echo-reply (0), echo-request (8), info-reply (16), info-request (15), mask-request (17), mask-reply (18), parameter-problem (12), redirect (5), router-advertisement (9), router-solicit (10), source-quench (4), time-exceeded (11), timestamp (13), timestamp-reply (14), or unreachable (3).</p>



NOTE: If you configure an interface with an input firewall filter that includes a reject action and with a service set that includes stateful firewall rules, the router executes the input firewall filter before the stateful firewall rules are run on the packet. As a result, when the Packet Forwarding Engine sends an ICMP error message out through the interface, the stateful firewall rules might drop the packet because it was not seen in the input direction.

Possible workarounds are to include a forwarding-table filter to perform the reject action, because this type of filter is executed after the stateful firewall in the input direction, or to include an output service filter to prevent the locally generated ICMP packets from going to the stateful firewall service.

Configuring Source and Destination Ports

The TCP or UDP source and destination port provide additional specification, in conjunction with the network protocol, for packet matching in an application definition. To configure ports, include the **destination-port** and **source-port** statements at the [edit applications application *application-name*] hierarchy level:

```
[edit applications application application-name]
destination-port value;
source-port value;
```

You must define one source or destination port. Normally, you specify this match in conjunction with the **protocol** match statement to determine which protocol is being used on the port; for constraints, see [Table 17 on page 304](#).

You can specify either a numeric value or one of the text synonyms listed in [Table 20 on page 309](#).

Table 20: Port Names Supported by Services Interfaces

Port Name	Corresponding Port Number
afs	1483
bgp	179
biff	512
bootpc	68
bootps	67
cmd	514
cvspserver	2401
dhcp	67

Table 20: Port Names Supported by Services Interfaces (*continued*)

Port Name	Corresponding Port Number
domain	53
eklogin	2105
ekshell	2106
exec	512
finger	79
ftp	21
ftp-data	20
http	80
https	443
ident	113
imap	143
kerberos-sec	88
klogin	543
kpasswd	761
krb-prop	754
krbupdate	760
kshell	544
ldap	389
login	513
mobileip-agent	434
mobileip-mn	435
msdp	639
netbios-dgm	138
netbios-ns	137

Table 20: Port Names Supported by Services Interfaces (*continued*)

Port Name	Corresponding Port Number
netbios-ssn	139
nfsd	2049
nntp	119
ntalk	518
ntp	123
pop3	110
pptp	1723
printer	515
radacct	1813
radius	1812
rip	520
rkinit	2108
smtp	25
snmp	161
snmptrap	162
snpp	444
socks	1080
ssh	22
sunrpc	111
syslog	514
tacacs-ds	65
talk	517
telnet	23
tftp	69

Table 20: Port Names Supported by Services Interfaces (*continued*)

Port Name	Corresponding Port Number
timed	525
who	513
xdmcp	177
zephyr-clt	2103
zephyr-hm	2104

For more information about matching criteria, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Configuring the Inactivity Timeout Period

You can specify a timeout period for application inactivity. If the software has not detected any activity during the duration, the flow becomes invalid when the timer expires. To configure a timeout period, include the `inactivity-timeout` statement at the `[edit applications application application-name]` hierarchy level:

```
[edit applications application application-name]
  inactivity-timeout seconds;
```

The default value is 30 seconds. The value you configure for an application overrides any global value configured at the `[edit interfaces interface-name service-options]` hierarchy level; for more information, see “[Configuring Default Timeout Settings for Services Interfaces](#)” on page 19.

Configuring SIP

The Session Initiation Protocol (SIP) is a generalized protocol for communication between endpoints involved in Internet services such as telephony, fax, video conferencing, instant messaging, and file exchange.

The Junos OS provides ALG services in accordance with the standard described in RFC 3261, *SIP: Session Initiation Protocol*. SIP flows under the Junos OS are as described in RFC 3665, *Session Initiation Protocol (SIP) Basic Call Flow Examples*.



NOTE: Before implementing the Junos OS SIP ALG, you should be familiar with certain limitations, discussed in “[Junos OS SIP ALG Limitations](#)” on page 319

The use of NAT in conjunction with the SIP ALG results in changes in SIP header fields due to address translation. For an explanation of these translations, refer to “[SIP ALG Interaction with Network Address Translation](#)” on page 313.

To implement SIP on adaptive services interfaces, you configure the **application-protocol** statement at the **[edit applications application *application-name*]** hierarchy level with the value **sip**. For more information about this statement, see [“Configuring an Application Protocol” on page 304](#). In addition, there are two other statements you can configure to modify how SIP is implemented:

- You can enable the router to accept any incoming SIP calls for the endpoint devices that are behind the NAT firewall. When a device behind the firewall registers with the proxy that is outside the firewall, the AS or Multiservices PIC maintains the registration state. When the **learn-sip-register** statement is enabled, the router can use this information to accept inbound calls. If this statement is not configured, no inbound calls are accepted; only the devices behind the firewall can call devices outside the firewall.

To configure SIP registration, include the **learn-sip-register** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]  
learn-sip-register;
```

You can also manually inspect the SIP register by issuing the **show services stateful-firewall sip-register** command; for more information, see the *Junos OS System Basics and Services Command Reference*.

- You can specify a timeout period for the duration of SIP calls that are placed on hold. When a call is put on hold, there is no activity and flows might time out after the configured **inactivity-timeout** period expires, resulting in call state teardown. To avoid this, when a call is put on hold, the flow timer is reset to the **sip-call-hold-timeout** cycle to preserve the call state and flows for longer than the **inactivity-timeout** period.

To configure a timeout period, include the **sip-call-hold-timeout** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]  
sip-call-hold-timeout seconds;
```

The default value is 7200 seconds and the range is from 0 through 36,000 seconds (10 hours).

SIP ALG Interaction with Network Address Translation

The Network Address Translation (NAT) protocol enables multiple hosts in a private subnet to share a single public IP address to access the Internet. For outgoing traffic, NAT replaces the private IP address of the host in the private subnet with the public IP address. For incoming traffic, the public IP address is converted back into the private address, and the message is routed to the appropriate host in the private subnet.

Using NAT with the Session Initiation Protocol (SIP) service is more complicated because SIP messages contain IP addresses in the SIP headers as well as in the SIP body. When using NAT with the SIP service, the SIP headers contain information about the caller and the receiver, and the device translates this information to hide it from the outside network. The SIP body contains the Session Description Protocol (SDP) information, which includes IP addresses and port numbers for transmission of the media. The device translates SDP information for allocating resources to send and receive the media.

How IP addresses and port numbers in SIP messages are replaced depends on the direction of the message. For an outgoing message, the private IP address and port number of the client are replaced with the public IP address and port number of the Juniper Networks firewall. For an incoming message, the public address of the firewall is replaced with the private address of the client.

When an INVITE message is sent out across the firewall, the SIP Application Layer Gateway (ALG) collects information from the message header into a call table, which it uses to forward subsequent messages to the correct endpoint. When a new message arrives, for example an ACK or 200 OK, the ALG compares the "From:", "To:", and "Call-ID:" fields against the call table to identify the call context of the message. If a new INVITE message arrives that matches the existing call, the ALG processes it as a REINVITE.

When a message containing SDP information arrives, the ALG allocates ports and creates a NAT mapping between them and the ports in the SDP. Because the SDP requires sequential ports for the Real-Time Transport Protocol (RTP) and Real-Time Control Protocol (RTCP) channels, the ALG provides consecutive even-odd ports. If it is unable to find a pair of ports, it discards the SIP message.

This topic contains the following sections:

- [Outgoing Calls on page 314](#)
- [Incoming Calls on page 315](#)
- [Forwarded Calls on page 315](#)
- [Call Termination on page 315](#)
- [Call Re-INVITE Messages on page 315](#)
- [Call Session Timers on page 316](#)
- [Call Cancellation on page 316](#)
- [Forking on page 316](#)
- [SIP Messages on page 316](#)
- [SIP Headers on page 316](#)
- [SIP Body on page 319](#)

Outgoing Calls

When a SIP call is initiated with a SIP request message from the internal to the external network, NAT replaces the IP addresses and port numbers in the SDP and binds the IP addresses and port numbers to the Juniper Networks firewall. Via, Contact, Route, and Record-Route SIP header fields, if present, are also bound to the firewall IP address. The ALG stores these mappings for use in retransmissions and for SIP response messages.

The SIP ALG then opens pinholes in the firewall to allow media through the device on the dynamically assigned ports negotiated based on information in the SDP and the Via, Contact, and Record-Route header fields. The pinholes also allow incoming packets to reach the Contact, Via, and Record-Route IP addresses and ports. When processing return traffic, the ALG inserts the original Contact, Via, Route, and Record-Route SIP fields back into packets.

Incoming Calls

Incoming calls are initiated from the public network to public static NAT addresses or to interface IP addresses on the device. Static NATs are statically configured IP addresses that point to internal hosts; interface IP addresses are dynamically recorded by the ALG as it monitors REGISTER messages sent by internal hosts to the SIP registrar. When the device receives an incoming SIP packet, it sets up a session and forwards the payload of the packet to the SIP ALG.

The ALG examines the SIP request message (initially an INVITE) and, based on information in the SDP, opens gates for outgoing media. When a 200 OK response message arrives, the SIP ALG performs NAT on the IP addresses and ports and opens pinholes in the outbound direction. (The opened gates have a short time-to-live, and they time out if a 200 OK response message is not received quickly.)

When a 200 OK response arrives, the SIP proxy examines the SDP information and reads the IP addresses and port numbers for each media session. The SIP ALG on the device performs NAT on the addresses and port numbers, opens pinholes for outbound traffic, and refreshes the timeout for gates in the inbound direction.

When the ACK arrives for the 200 OK, it also passes through the SIP ALG. If the message contains SDP information, the SIP ALG ensures that the IP addresses and port numbers are not changed from the previous INVITE—if they are, the ALG deletes old pinholes and creates new pinholes to allow media to pass through. The ALG also monitors the Via, Contact, and Record-Route SIP fields and opens new pinholes if it determines that these fields have changed.

Forwarded Calls

A forwarded call is when, for example, user A outside the network calls user B inside the network, and user B forwards the call to user C outside the network. The SIP ALG processes the INVITE from user A as a normal incoming call. But when the ALG examines the forwarded call from B to C outside the network and notices that B and C are reached using the same interface, it does not open pinholes in the firewall, because media will flow directly between user A and user C.

Call Termination

The BYE message terminates a call. When the device receives a BYE message, it translates the header fields just as it does for any other message. But because a BYE message must be acknowledged by the receiver with a 200 OK, the ALG delays call teardown for five seconds to allow time for transmission of the 200 OK.

Call Re-INVITE Messages

Re-INVITE messages add new media sessions to a call and remove existing media sessions. When new media sessions are added to a call, new pinholes are opened in the firewall and new address bindings are created. The process is identical to the original call setup. When one or more media sessions are removed from a call, pinholes are closed and bindings released just as with a BYE message.

Call Session Timers

The SIP ALG uses the Session-Expires value to time out a session if a Re-INVITE or UPDATE message is not received. The ALG gets the Session-Expires value, if present, from the 200 OK response to the INVITE and uses this value for signaling timeout. If the ALG receives another INVITE before the session times out, it resets all timeout values to this new INVITE or to default values, and the process is repeated.

As a precautionary measure, the SIP ALG uses hard timeout values to set the maximum amount of time a call can exist. This ensures that the device is protected should one of the following events occur:

- End systems crash during a call and a BYE message is not received.
- Malicious users never send a BYE in an attempt to attack a SIP ALG.
- Poor implementations of SIP proxy fail to process Record-Route and never send a BYE message.
- Network failures prevent a BYE message from being received.

Call Cancellation

Either party can cancel a call by sending a CANCEL message. Upon receiving a CANCEL message, the SIP ALG closes pinholes through the firewall—if any have been opened—and releases address bindings. Before releasing the resources, the ALG delays the control channel age-out for approximately five seconds to allow time for the final 200 OK to pass through. The call is terminated when the five second timeout expires, regardless of whether a 487 or non-200 response arrives.

Forking

Forking enables a SIP proxy to send a single INVITE message to multiple destinations simultaneously. When the multiple 200 OK response messages arrive for the single call, the SIP ALG parses but updates call information with the first 200 OK messages it receives.

SIP Messages

The SIP message format consists of a SIP header section and the SIP body. In request messages, the first line of the header section is the request line, which includes the method type, request-URI, and protocol version. In response messages, the first line is the status line, which contains a status code. SIP headers contain IP addresses and port numbers used for signaling. The SIP body, separated from the header section by a blank line, is reserved for session description information, which is optional. Junos OS currently supports the SDP only. The SIP body contains IP addresses and port numbers used to transport the media.

SIP Headers

In the following sample SIP request message, NAT replaces the IP addresses in the header fields to hide them from the outside network.

```
INVITE bob@10.150.20.5 SIP/2.0
Via: SIP/2.0/UDP 10.150.20.3:5434
```

From: alice@10.150.20.3
To: bob@10.150.20.5
Call-ID: a12abcde@10.150.20.3
Contact: alice@10.150.20.3:5434
Route: <sip:netscreen@10.150.20.3:5060>
Record-Route: <sip:netscreen@10.150.20.3:5060>

How IP address translation is performed depends on the type and direction of the message. A message can be any of the following:

- Inbound request
- Outbound response
- Outbound request
- Inbound response

Table 21 on page 317 shows how NAT is performed in each of these cases. Note that for several of the header fields the ALG determine more than just whether the messages comes from inside or outside the network. It must also determine what client initiated the call, and whether the message is a request or response.

Table 21: Requesting Messages with NAT Table

Inbound Request (from public to private)	To:	Replace domain with local address
	From:	None
	Call-ID:	None
	Via:	None
	Request-URI:	Replace ALG address with local address
	Contact:	None
	Record-Route:	None
	Route:	None

Table 21: Requesting Messages with NAT Table (*continued*)

Outbound Response (from private to public)	To:	Replace ALG address with local address
	From:	None
	Call-ID:	None
	Via:	None
	Request-URI:	N/A
	Contact:	Replace local address with ALG address
	Record-Route:	Replace local address with ALG address
	Route:	None
Outbound Request (from private to public)	To:	None
	From:	Replace local address with ALG address
	Call-ID:	None
	Via:	Replace local address with ALG address
	Request-URI:	None
	Contact:	Replace local address with ALG address
	Record-Route:	Replace local address with ALG address
	Route:	Replace ALG address with local address
Outbound Response (from public to private)	To:	None
	From:	Replace ALG address with local address
	Call-ID:	None
	Via:	Replace ALG address with local address
	Request-URI:	N/A
	Contact:	None
	Record-Route:	Replace ALG address with local address
	Route:	Replace ALG address with local address

SIP Body

The SDP information in the SIP body includes IP addresses the ALG uses to create channels for the media stream. Translation of the SDP section also allocates resources, that is, port numbers to send and receive the media.

The following excerpt from a sample SDP section shows the fields that are translated for resource allocation.

```
o=user 2344234 55234434 IN IP4 10.150.20.3
c=IN IP4 10.150.20.3
m=audio 43249 RTP/AVP 0
```

SIP messages can contain more than one media stream. The concept is similar to attaching multiple files to an e-mail message. For example, an INVITE message sent from a SIP client to a SIP server might have the following fields:

```
c=IN IP4 10.123.33.4
m=audio 33445 RTP/AVP 0
c=IN IP4 10.123.33.4
m=audio 33447 RTP/AVP 0
c=IN IP4 10.123.33.4
m=audio 33449 RTP/AVP 0
```

Junos OS supports up to 6 SDP channels negotiated for each direction, for a total of 12 channels per call. For more information, see *SDP Session Descriptions*.

Junos OS SIP ALG Limitations

The following limitations apply to configuration of the SIP ALG:

- Only the methods described in RFC 3261 are supported.
- Only SIP version 2 is supported.
- TCP is not supported as a transport mechanism for signaling messages.
- *Do not configure the SIP ALG when using STUN.* If clients use STUN/TURN to detect the firewall or NAT devices between the caller and responder or proxy, the client attempts to best-guess the NAT device behavior and act accordingly to place the call.
- Do not use the endpoint-independent mapping NAT pool option in conjunction with the SIP ALG. Errors will result.
- IPv6 signaling data is not supported.
- Authentication is not supported.
- Encrypted messages are not supported.
- SIP fragmentation is not supported.
- The maximum UDP packet size containing a SIP message is assumed to be 4 KB. SIP messages larger than this are not supported.
- The maximum number of media channels in a SIP message is assumed to be six.
- Fully qualified domain names (FQDNs) are not supported in critical fields.

- QoS is not supported.
- High availability is not supported, except for warm standby.
- A timeout setting of never is not supported on SIP or NAT.
- Multicast (forking proxy) is not supported.

Configuring an SNMP Command for Packet Matching

You can specify an SNMP command setting for packet matching. To configure SNMP, include the **snmp-command** statement at the **[edit applications application application-name]** hierarchy level:

```
[edit applications application application-name]
snmp-command value;
```

The supported values are **get**, **get-next**, **set**, and **trap**. You can configure only one value for matching. The **application-protocol** statement at the **[edit applications application application-name]** hierarchy level must have the value **snmp**. For information about specifying the application protocol, see [“Configuring an Application Protocol” on page 304](#).

Configuring an RPC Program Number

You can specify an RPC program number for packet matching. To configure an RPC program number, include the **rpc-program-number** statement at the **[edit applications application application-name]** hierarchy level:

```
[edit applications application application-name]
rpc-program-number number;
```

The range of values used for DCE or RPC is from 100,000 through 400,000. The **application-protocol** statement at the **[edit applications application application-name]** hierarchy level must have the value **rpc**. For information about specifying the application protocol, see [“Configuring an Application Protocol” on page 304](#).

Configuring the TTL Threshold

You can specify a trace route time-to-live (TTL) threshold value, which controls the acceptable level of network penetration for trace routing. To configure a TTL value, include the **ttl-threshold** statement at the **[edit applications application application-name]** hierarchy level:

```
[edit applications application application-name]
ttl-threshold value;
```

The **application-protocol** statement at the **[edit applications application application-name]** hierarchy level must have the value **traceroute**. For information about specifying the application protocol, see [“Configuring an Application Protocol” on page 304](#).

Configuring a Universal Unique Identifier

You can specify a Universal Unique Identifier (UUID) for DCE RPC objects. To configure a UUID value, include the **uuid** statement at the **[edit applications application application-name]** hierarchy level:

```
[edit applications application application-name]
  uuid hex-value;
```

The **uuid** value is in hexadecimal notation. The **application-protocol** statement at the **[edit applications **application** *application-name*** hierarchy level must have the value **dce-rpc**. For information about specifying the application protocol, see “[Configuring an Application Protocol](#)” on page 304. For more information on UUID numbers, see <http://www.opengroup.org/onlinepubs/9629399/apdxa.htm>.

Examples: Configuring Application Protocols

The following example shows an application protocol definition describing a special FTP application running on port 78:

```
[edit applications]
application my-ftp-app {
  application-protocol ftp;
  protocol tcp;
  destination-port 78;
  timeout 100; # inactivity timeout for FTP service
}
```

The following example shows a special ICMP protocol (**application-protocol icmp**) of type 8 (ICMP echo):

```
[edit applications]
application icmp-app {
  application-protocol icmp;
  protocol icmp;
  icmp-type icmp-echo;
}
```

The following example shows a possible application set:

```
[edit applications]
application-set basic {
  http;
  ftp;
  telnet;
  nfs;
  icmp;
}
```

The software includes a predefined set of well-known application protocols. The set includes applications for which the TCP and UDP destination ports are already recognized by stateless firewall filters.

Related Documentation

- [ALG Descriptions on page 277](#)
- [Configuring Application Sets on page 303](#)
- [Configuring Application Protocol Properties on page 303](#)
- [Verifying the Output of ALG Sessions](#)

ICMP, Ping, and Traceroute ALGs for MS-MICs and MS-MPCs

Starting with Junos OS Release 14.2, Junos OS extension-provider packages that are preinstalled and preconfigured on the MS-MIC and MS-MPC offer support for ping, traceroute, and ICMP ALGs in a consistent manner that is identical to the support that the uKernel service provides. Parity and uniformity of support is established for these ALGs between MS-MICs/MS-MPCs and the uKernel service. Until Junos OS Release 14.1, ICMP ALGs, ping ALGs, and traceroute ALGs were not entirely supported on MX Series routers with MS-MICs and MS-MPCs in comparison with the uKernel service that enables Network Address Translation (NAT) with stateful firewall (SFW) on the uKernel PIC. Support was available for handling of ICMP error response packets that match any existing flow in the opposite direction and NAT processing of ICMP packets with NAT processing of ping packets.

On MX Series routers with MS-MICs and MS-MPCs, tracking of ping traffic states wholly using the ICMP sequence numbers (for example, forwarding an echo reply only if the echo request with the corresponding sequence number is identified) is supported. ICMP application layer gateway (ALG) is enhanced to provide detailed logging information. Also, the traceroute ALGs enable UDP probe packets to be processed with the UDP destination port number greater than 33000 and the IP time-to-live (TTL) is less than 30 seconds. Traceroute ALGs enable ICMP response packets for which the ICMP type is time-exceeded to be processed and support a traceroute TTL threshold value, which controls the acceptable level of network penetration for trace routing.

You can configure ICMP and ping messages with the **application junos-icmp-all**, **application junos-icmp-ping**, and **application icmp-code** statements at the **[edit services stateful-firewall rule rule-name term term-name from]** and the **[edit services nat rule rule-name term term-name from]** hierarchy levels to define the match condition for the stateful firewall and NAT rules. Until Junos OS Release 14.1, a restriction or a validation on the applications that you could define for ICMP messages was not present. MS-MICs and MS-MPCs function the same way as the uKernel service, which causes the ping traffic to be tracked statefully using the ICMP sequence numbers (an echo reply is forwarded only if echo request with the corresponding sequence number matches). Also, MS-MICs and MS-MPCs impose a limit on the outstanding ping requests and drop the subsequent ping requests when the limit is reached.

Similarly, for traceroute messages, you can configure the **application junos-traceroute** and **application junos-traceroute-ttl-1** statements at the **[edit services stateful-firewall rule rule-name term term-name from]** and the **[edit services nat rule rule-name term term-name from]** hierarchy levels to define the match condition for traceroute messages for the stateful firewall and NAT rules.

Traceroute and ICMP messages are supported for both IPv4 and IPv6 packets. For the traceroute functionality to work, you only need to ensure that the user-defined applications are working as expected with the inactivity timeout period and the TTL threshold values are configured to be the same period of time as configured by using the **session-timeout seconds** statement at the **[edit services application-identification application application-name]** hierarchy level. During the logging of ICMP messages, the ALG information for ping and ICMP utilities is displayed in the output of the relevant show

commands, such as **show sessions** and **show conversations**, in the same manner as displayed for uKernel logging.

Related Documentation

- [ALG Descriptions on page 277](#)

Monitoring Port Control Protocol Operations

You can monitor Port Control Protocol (PCP) operations with the following operational commands:

- **show services nat mappings pcp**
- **show services nat mappings endpoint-independent**
- **show services pcp statistics protocol**

The following are examples of the output of these commands.

```
user@host> show services nat mappings pcp
```

```
Interface: sp-0/0/0, Service set: in
```

```
NAT pool: p
PCP Client      : 10.1.1.2          PCP lifetime : 995
Mapping         : 10.1.1.2          : 9000 --> 8.8.8.8      : 1025
Session Count   : 1
Mapping State    : Active
```

```
DS-LITE output:
```

```
=====
PCP Client      : 2222::1          PCP lifetime : 106
Mapping         : 88.1.0.47        : 47 --> 70.70.70.1    :41972
Session Count   : 1
Mapping State    : Active
B4 Address      : 2222::1
```

```
user@host> show services nat mappings endpoint-independent
```

```
Interface: sp-0/0/0, Service set: in
```

```
NAT pool: p
Mapping         : 10.1.1.2          :57400 --> 8.8.8.8      : 1024
Session Count   : 0
Mapping State    : Timeout
PCP Client      : 10.1.1.2          PCP lifetime : 991
Mapping         : 10.1.1.2          : 9000 --> 8.8.8.8      : 1025
Session Count   : 1
Mapping State    : Active
```

```
DS-LITE output:
```

```
=====
PCP Client      : 2222::1          PCP lifetime : 190
Mapping         : 88.1.1.3          : 4001 --> 70.70.70.2    :58989
Session Count   : 1
Mapping State    : Active
B4 Address      : 2222::1
```

```
user@host> show services pcp statistics protocol
```

```
Protocol Statistics:
```

Operational Statistics

Map request received	:0
Peer request received	:0
Other operational counters	:0

Option Statistics

Unprocessed requests received	:0
Third party requests received	:0
Prefer fail option received	:0
Filter option received	:0
Other options counters	:0
Option optional received	:0

Result Statistics

PCP success	:0
PCP unsupported version	:0
Not authorized	:0
Bad requests	:0
Unsupported opcode	:0
Unsupported option	:0
Bad option	:0
Network failure	:0
Out of resources	:0
Unsupported protocol	:0
User exceeded quota	:0
Cannot provide external	:0
Address mismatch	:0
Excessive number of remote peers	:0
Processing error	:0
Other result counters	:0

PART 6

Securing Content Using Junos Network Secure and IDS

- [Junos Network Secure Overview on page 327](#)
- [Junos Network Secure Configuration Overview on page 331](#)
- [IDS Configuration Overview on page 355](#)
- [Monitoring Junos Network Secure on page 367](#)

Junos Network Secure Overview

- [Junos Network Secure Overview on page 327](#)

Junos Network Secure Overview

Routers use firewalls to track and control the flow of traffic. The following platforms employ a type of firewall called a *stateful firewall*.

- MultiServices Dense Port Concentrators (MS-DPCs)
- MS-100, MS-400, and MS-500 MultiServices PICs
- MultiServices Modular Port Concentrators (MS-MPCs), and Multiservices Modular Interface Cards (MS-MICs)

The stateful firewall capabilities provided by the Junos OS are collectively known as *Junos Network Secure*.

Contrasted with a *stateless* firewall that inspects packets in isolation, a stateful firewall provides an extra layer of security by using state information derived from past communications and other applications to make dynamic control decisions for new communication attempts.

Stateful firewalls group relevant *flows* into *conversations*. A flow is identified by the following five properties:

- Source address
- Source port
- Destination address
- Destination port
- Protocol

A typical Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) conversation consists of two flows: the initiation flow and the responder flow. However, some conversations, such as an FTP conversation, might consist of two control flows and many data flows.

Firewall rules govern whether the conversation is allowed to be established. If a conversation is allowed, all flows within the conversation are permitted, including flows that are created during the life cycle of the conversation.

You configure stateful firewalls using a powerful rule-driven conversation handling path. A *rule* consists of direction, source address, source port, destination address, destination port, IP protocol value, and application protocol or service. In addition to the specific values you configure, you can assign the value **any** to rule objects, addresses, or ports, which allows them to match any input value. Finally, you can optionally negate the rule objects, which negates the result of the type-specific match.

Firewall rules are directional. For each new conversation, the router software checks the initiation flow matching the direction specified by the rule.

Firewall rules are ordered. The software checks the rules in the order in which you include them in the configuration. The first time the firewall discovers a match, the router implements the action specified by that rule. Rules still unchecked are ignored.



NOTE: Starting with Junos OS Release 14.2, MS-MPC and MS-MIC interface cards support IPv6 traffic for Junos Network Secure Stateful Firewall.

For more information, see [“Configuring Stateful Firewall Rules” on page 331](#).

Stateful Firewall Support for Application Protocols

By inspecting the application protocol data, the Junos Network Secure stateful firewall can intelligently enforce security policies and allow only the minimal required packet traffic to flow through the firewall.

The firewall rules are configured in relation to an interface. By default, the stateful firewall allows all sessions initiated from the hosts behind the interface to pass through the router.

Stateful Firewall Anomaly Checking

The stateful firewall recognizes the following events as anomalies and sends them to the IDS software for processing:

- IP anomalies:
 - IP version is not correct.
 - IP header length field is too small.
 - IP header length is set larger than the entire packet.
 - Bad header checksum.
 - IP total length field is shorter than header length.
 - Packet has incorrect IP options.

- Internet Control Message Protocol (ICMP) packet length error.
- Time-to-live (TTL) equals 0.
- IP address anomalies:
 - IP packet source is a broadcast or multicast.
 - Land attack (source IP equals destination IP).
- IP fragmentation anomalies:
 - IP fragment overlap.
 - IP fragment missed.
 - IP fragment length error.
 - IP packet length is more than 64 kilobytes (KB).
 - Tiny fragment attack.
- TCP anomalies:
 - TCP port 0.
 - TCP sequence number 0 and flags 0.
 - TCP sequence number 0 and FIN/PSH/RST flags set.
 - TCP flags with wrong combination (TCP FIN/RST or SYN/(URG|FIN|RST)).
 - Bad TCP checksum.
- UDP anomalies:
 - UDP source or destination port 0.
 - UDP header length check failed.
 - Bad UDP checksum.
- Anomalies found through stateful TCP or UDP checks:
 - SYN followed by SYN-ACK packets without ACK from initiator.
 - SYN followed by RST packets.
 - SYN without SYN-ACK.
 - Non-SYN first flow packet.
 - ICMP unreachable errors for SYN packets.
 - ICMP unreachable errors for UDP packets.
- Packets dropped according to stateful firewall rules.

If you employ stateful anomaly detection in conjunction with stateless detection, IDS can provide early warning for a wide range of attacks, including these:

- TCP or UDP network probes and port scanning

- SYN flood attacks
- IP fragmentation-based attacks such as teardrop, bonk, and boink

Junos Network Secure Configuration Overview

- [Configuring Stateful Firewall Rules on page 331](#)
- [Configuring Stateful Firewall Rule Sets on page 335](#)
- [Examples: Configuring Stateful Firewall Rules on page 335](#)
- [Example: BOOTP and Broadcast Addresses on page 338](#)
- [Example: Configuring Layer 3 Services and the Services SDK on Two PICs on page 339](#)
- [Example: Virtual Routing and Forwarding \(VRF\) and Service Configuration on page 351](#)

Configuring Stateful Firewall Rules

To configure a stateful firewall rule, include the **rule** *rule-name* statement at the **[edit services stateful-firewall]** hierarchy level:

```
[edit services stateful-firewall]
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application-sets set-name;
      applications [ application-names ];
      destination-address address <except>;
      destination-address-range low minimum-value high maximum-value <except>;
      destination-prefix-list list-name <except>;
      source-address address <except>;
      source-address-range low minimum-value high maximum-value <except>;
      source-prefix-list list-name <except>;
    }
    then {
      (accept | discard | reject);
      allow-ip-options [ values ];
      syslog;
    }
  }
}
```

Each stateful firewall rule consists of a set of terms, similar to a filter configured at the **[edit firewall]** hierarchy level. A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded. The **from** statement is optional in stateful firewall rules.
- **then** statement—Specifies the actions and action modifiers to be performed by the router software. The **then** statement is mandatory in stateful firewall rules.

The following sections explain how to configure the components of stateful firewall rules:

- [Configuring Match Direction for Stateful Firewall Rules on page 332](#)
- [Configuring Match Conditions in Stateful Firewall Rules on page 332](#)
- [Configuring Actions in Stateful Firewall Rules on page 334](#)

Configuring Match Direction for Stateful Firewall Rules

Each rule must include a **match-direction** statement that specifies the direction in which the rule match is applied. To configure where the match is applied, include the **match-direction** statement at the **[edit services stateful-firewall rule rule-name]** hierarchy level:

```
[edit services stateful-firewall rule rule-name]
match-direction (input | output | input-output);
```

If you configure **match-direction input-output**, sessions initiated from both directions might match this rule.

The match direction is used with respect to the traffic flow through the AS or Multiservices PIC. When a packet is sent to the PIC, direction information is carried along with it.

With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.

With a next-hop service set, packet direction is determined by the interface used to route the packet to the AS or Multiservices PIC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the PIC, the packet direction is output. For more information on inside and outside interfaces, see [“Configuring Service Sets to be Applied to Services Interfaces” on page 31](#).

On the PIC, a flow lookup is performed. If no flow is found, rule processing is performed. Rules in this service set are considered in sequence until a match is found. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that matches the packet direction are considered. Most packets result in the creation of bidirectional flows.

Configuring Match Conditions in Stateful Firewall Rules

To configure stateful firewall match conditions, include the **from** statement at the **[edit services stateful-firewall rule rule-name term term-name]** hierarchy level:

```
[edit services stateful-firewall rule rule-name term term-name]
```

```

from {
  application-sets set-name;
  applications [ application-names ];
  destination-address (address | any-unicast) <except>;
  destination-address-range low minimum-value high maximum-value <except>;
  destination-prefix-list list-name <except>;
  source-address (address | any-unicast) <except>;
  source-address-range low minimum-value high maximum-value <except>;
  source-prefix-list list-name <except>;
}

```

The source address and destination address can be either IPv4 or IPv6. You can use either the source address or the destination address as a match condition, in the same way that you would configure a firewall filter; for more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*. You can use the wildcard value **any-unicast**, which denotes matching all unicast addresses.

Alternatively, you can specify a list of source or destination prefixes by configuring the **prefix-list** statement at the **[edit policy-options]** hierarchy level and then including either the **destination-prefix-list** or the **source-prefix-list** statement in the stateful firewall rule. For an example, see “[Examples: Configuring Stateful Firewall Rules](#)” on page 335.

If you omit the **from** term, the stateful firewall accepts all traffic and the default protocol handlers take effect:

- User Datagram Protocol (UDP), Transmission Control Protocol (TCP), and Internet Control Message Protocol (ICMP) create a bidirectional flow with a predicted reverse flow.
- IP creates a unidirectional flow.

You can also include application protocol definitions you have configured at the **[edit applications]** hierarchy level; for more information, see “[Configuring Application Protocol Properties](#)” on page 303.

- To apply one or more specific application protocol definitions, include the **applications** statement at the **[edit services stateful-firewall rule rule-name term term-name from]** hierarchy level.
- To apply one or more sets of application protocol definitions you have defined, include the **application-sets** statement at the **[edit services stateful-firewall rule rule-name term term-name from]** hierarchy level.



NOTE: If you include one of the statements that specifies application protocols, the router derives port and protocol information from the corresponding configuration at the **[edit applications]** hierarchy level; you cannot specify these properties as match conditions.

Configuring Actions in Stateful Firewall Rules

To configure stateful firewall actions, include the **then** statement at the **[edit services stateful-firewall rule *rule-name* term *term-name*]** hierarchy level:

```
[edit services stateful-firewall rule rule-name term term-name]  
then {  
  (accept | discard | reject);  
  allow-ip-options [ values ];  
  syslog;  
}
```

You must include one of the following three possible actions:

- **accept**—The packet is accepted and sent on to its destination.
- **discard**—The packet is not accepted and is not processed further.
- **reject**—The packet is not accepted and a rejection message is returned; UDP sends an ICMP unreachable code and TCP sends RST. Rejected packets can be logged or sampled.

You can optionally configure the firewall to record information in the system logging facility by including the **syslog** statement at the **[edit services stateful-firewall rule *rule-name* term *term-name* then]** hierarchy level. This statement overrides any **syslog** setting included in the service set or interface default configuration.

Configuring IP Option Handling

You can optionally configure the firewall to inspect IP header information by including the **allow-ip-options** statement at the **[edit services stateful-firewall rule *rule-name* term *term-name* then]** hierarchy level. When you configure this statement, all packets that match the criteria specified in the **from** statement are subjected to additional matching criteria. A packet is accepted only when all of its IP option types are configured as values in the **allow-ip-options** statement. If you do not configure **allow-ip-options**, only packets without IP header options are accepted.

The additional IP header option inspection applies only to the **accept** and **reject** stateful firewall actions. This configuration has no effect on the **discard** action. When the IP header inspection fails, reject frames are not sent; in this case, the **reject** action has the same effect as **discard**.

If an IP option packet is accepted by the stateful firewall, Network Address Translation (NAT) and intrusion detection service (IDS) are applied in the same way as to packets without IP option headers. The IP option configuration appears only in the stateful firewall rules; NAT applies to packets with or without IP options.

When a packet is dropped because it fails the IP option inspection, this exception event generates both IDS event and system log messages. The event type depends on the first IP option field rejected.

[Table 22 on page 335](#) lists the possible values for the **allow-ip-options** statement. You can include a range or set of numeric values, or one or more of the predefined IP option

settings. You can enter either the option name or its numeric equivalent. For more information, refer to <http://www.iana.org/assignments/ip-parameters>.

Table 22: IP Option Values

IP Option Name	Numeric Value	Comment
any	0	Any IP option
ip-security	130	–
ip-stream	136	–
loose-source-route	131	–
route-record	7	–
router-alert	148	–
strict-source-route	137	–
timestamp	68	–

Configuring Stateful Firewall Rule Sets

The **rule-set** statement defines a collection of stateful firewall rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services stateful-firewall]** hierarchy level with a **rule** statement for each rule:

```
[edit services stateful-firewall]
rule-set rule-set-name {
  rule rule-name;
}
```

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

Examples: Configuring Stateful Firewall Rules

The following example show a stateful firewall configuration containing two rules, one for input matching on a specified application set and the other for output matching on a specified source address:

```
[edit services]
stateful-firewall {
```

```
rule Rule1 {
  match-direction input;
  term 1 {
    from {
      application-sets Applications;
    }
    then {
      accept;
    }
  }
  term accept {
    then {
      accept;
    }
  }
}
rule Rule2 {
  match-direction output;
  term Local {
    from {
      source-address {
        10.1.3.2/32;
      }
    }
    then {
      accept;
    }
  }
}
```

The following example has a single rule with two terms. The first term rejects all traffic in **my-application-group** that originates from the specified source address, and provides a detailed system log record of the rejected packets. The second term accepts Hypertext Transfer Protocol (HTTP) traffic from anyone to the specified destination address.

```
[edit services stateful-firewall]
rule my-firewall-rule {
  match-direction input-output;
  term term1 {
    from {
      source-address 10.1.3.2/32;
      application-sets my-application-group;
    }
    then {
      reject;
      syslog;
    }
  }
  term term2 {
    from {
      destination-address 10.2.3.2/32;
      applications http;
    }
    then {
      accept;
    }
  }
}
```

```

    }
  }
}

```

The following example shows use of source and destination prefix lists. This requires two separate configuration items.

You configure the prefix list at the **[edit policy-options]** hierarchy level:

```

[edit]
policy-options {
  prefix-list p1 {
    1.1.1.1/32;
    2.2.2.0/24;
  }
  prefix-list p2 {
    3.3.3.3/32;
    4.4.4.0/24;
  }
}

```

You reference the configured prefix list in the stateful firewall rule:

```

[edit]
services {
  stateful-firewall {
    rule r1 {
      match-direction input;
      term t1 {
        from {
          source-prefix-list {
            p1;
          }
          destination-prefix-list {
            p2;
          }
        }
        then {
          accept;
        }
      }
    }
  }
}

```

This is equivalent to the following configuration:

```

[edit]
services {
  stateful-firewall {
    rule r1 {
      match-direction input;
      term t1 {
        from {
          source-address {
            1.1.1.1/32;
            2.2.2.0/24;
          }
        }
      }
    }
  }
}

```

```
    }
    destination-address {
        3.3.3.3/32;
        4.4.4.0/24;
    }
}
then {
    accept;
}
}
}
```

You can use the **except** qualifier with the prefix lists, as in the following example. In this case, the **except** qualifier applies to all prefixes included in prefix list **p2**.

```
[edit]
services {
    stateful-firewall {
        rule r1 {
            match-direction input;
            term t1 {
                from {
                    source-prefix-list {
                        p1;
                    }
                    destination-prefix-list {
                        p2 except;
                    }
                }
                then {
                    accept;
                }
            }
        }
    }
}
```

For additional examples that combine stateful firewall configuration with other services and with virtual private network (VPN) routing and forwarding (VRF) tables, see the configuration examples.

Related Documentation

- [Example: BOOTP and Broadcast Addresses on page 338](#)
- [Example: Dynamic Source NAT as a Next-Hop Service on page 138](#)
- [Example: Virtual Routing and Forwarding \(VRF\) and Service Configuration on page 351](#)
- [Example: Service Interfaces Configuration on page 16](#)
- [Example: Configuring Layer 3 Services and the Services SDK on Two PICs on page 339](#)

Example: BOOTP and Broadcast Addresses

The following example supports Bootstrap Protocol (BOOTP) and broadcast addresses:

```

[edit applications]
application bootp {
  application-protocol bootp;
  protocol udp;
  destination-port 67;
}
[edit services]
stateful-firewall bootp-support {
  rule bootp-allow {
    direction input;
    term bootp-allow {
      from {
        destination-address {
          any-unicast;
          255.255.255.255;
        }
        application bootp;
      }
      then {
        accept;
      }
    }
  }
}
}

```

Example: Configuring Layer 3 Services and the Services SDK on Two PICs

You can configure the Layer 3 service package and the Services SDK on two PICs. For this example, you must configure an FTP or HTTP client and a server. In this configuration, the client side of the router interface is ge-1/2/2.1 and the server side of the router interface is ge-1/1/0.48. This configuration enables Network Address Translation (NAT) with stateful firewall (SFW) on the uKernel PIC and application identification (APPID), application-aware access list (AACL), and intrusion detection and prevention (IDP) on the Services SDK PIC for FTP or HTTP traffic.



NOTE: The Services SDK does not support NAT yet. When NAT is required, you can configure the Layer 3 service package to deploy NAT along with the Services SDK such as APPID, AACL, or IDP.

To deploy the Layer 3 service package and the Services SDK on two PICs:

1. In configuration mode, go to the following hierarchy level:

```

[edit services]
user@host# edit stateful-firewall

```

2. In the hierarchy level, configure the conditions for the stateful firewall rule r1.

```

[edit services stateful-firewall]
user@host# set rule rule-name match-direction input-output term term from
  applications application-name
user@host# set rule rule-name match-direction input-output term term then accept
  syslog

```

In this example, the stateful firewall term is **ALLOWED-SERVICES**. Enclose the application names—`junos-ftp`, `junos-http`, and `junos-icmp-ping`—in brackets for *application-name*.

```
[edit services stateful-firewall]
user@host# set rule r1 match-direction input-output term ALLOWED-SERVICES from
  applications [ junos-ftp junos-http junos-icmp-ping ]
user@host# set rule r1 match-direction input-output term ALLOWED-SERVICES then
  accept syslog
```

3. Configure the conditions for the stateful firewall rule `r2`.

```
[edit services stateful-firewall]
user@host# set rule rule-name match-direction input-output term term then discard
user@host# set rule rule-name match-direction input-output term term then syslog
```

In this example, the stateful firewall term is `term1`.

```
[edit services stateful-firewall]
user@host# set rule r2 match-direction input-output term term1 then discard
user@host# set rule r2 match-direction input-output term term1 then syslog
```

4. Go to the following hierarchy level and verify the configuration:

```
[edit services stateful-firewall]
user@host# show
rule r1 {
  match-direction input-output;
  term ALLOWED-SERVICES {
    from {
      applications [ junos-ftp junos-http junos-icmp-ping ];
    }
    then {
      accept;
      syslog;
    }
  }
}
rule r2 {
  match-direction input-output;
  term term1 {
    then {
      discard;
      syslog;
    }
  }
}
```

5. Go to the following hierarchy level:

```
[edit services]
user@host# edit nat
```

6. In the hierarchy level, configure the NAT pool.

```
[edit services nat]
user@host# set pool pool-name address ip-address
user@host# set pool pool-name port automatic
```

In this example, the NAT pool is **OUTBOUND-SERVICES** and the IP address is **10.48.0.2/32**.

```
[edit services nat]
user@host# set pool OUTBOUND-SERVICES address 10.48.0.2/32
user@host# set pool OUTBOUND-SERVICES port automatic
```

7. Configure the NAT rule.

```
[edit services nat]
user@host# set rule rule-name match-direction output term term from applications
application-name
user@host# set rule rule-name match-direction output term term then translated
source-pool source-pool translation-type source dynamic
```

In this example, the NAT rule is **SET-MSR-ADDR**, the NAT term is **TRANSLATE-SOURCE-ADDR**, and the source pool is **OUTBOUND-SERVICES**. Enclose the application names—**junos-ftp**, **junos-http**, and **junos-icmp-ping**—in parentheses for *application-name*.

```
[edit services nat]
user@host# set rule SET-MSR-ADDR match-direction output term
TRANSLATE-SOURCE-ADDR from applications [ junos-ftp junos-http
junos-icmp-ping ]
user@host# set rule SET-MSR-ADDR match-direction output term
TRANSLATE-SOURCE-ADDR then translated source-pool OUTBOUND-SERVICES
translation-type source dynamic
```

8. Go to the following hierarchy level and verify the configuration:

```
[edit services nat]
user@host# show
pool OUTBOUND-SERVICES {
    address 11.48.0.2/32;
    port {
        automatic;
    }
}
rule SET-MSR-ADDR {
    match-direction output;
    term TRANSLATE-SOURCE-ADDR {
        from {
            applications [ junos-ftp junos-http junos-icmp-ping ];
        }
        then {
            translated {
                source-pool OUTBOUND-SERVICES;
                translation-type {
                    source dynamic;
                }
            }
        }
    }
}
```

9. Go to the following hierarchy level:

```
[edit security]
user@host# edit idp
```

10. In the hierarchy level, configure the IDP policy.

```
[edit security idp]
user@host# set idp-policy policy-name rulebase-ips rule rule-name match application
    default attacks predefined-attacks attack-name
user@host# set idp-policy policy-name rulebase-ips rule rule-name match application
    default attacks predefined-attack-groups attack-group--name
user@host# set idp-policy policy-name rulebase-ips rule rule-name then action
    no-action
user@host# set idp-policy policy-name rulebase-ips rule rule-name then notification
    log-attacks alert
```

In this example, the IDP policy is **test1**, the rule is **r1**, the predefined attack is **FTP:USER:ROOT**, and the predefined attack group is **"Recommended Attacks"**.

```
[edit security idp]
user@host# set idp-policy test1 rulebase-ips rule r1 match application default attacks
    predefined-attacks FTP:USER:ROOT
user@host# set idp-policy test1 rulebase-ips rule r1 match application default attacks
    predefined-attack-groups [ "Recommended Attacks" ]
user@host# set idp-policy test1 rulebase-ips rule r1 then action no-action
user@host# set idp-policy test1 rulebase-ips rule r1 then notification log-attacks alert
```

11. Configure the trace options for IDP services.

```
[edit security idp]
user@host# set traceoptions file filename
user@host# set traceoptions flag all
user@host# set traceoptions level all
```

In this example, the log file name is **idp-demo.log**.

```
[edit security idp]
user@host# set traceoptions file idp-demo.log
user@host# set traceoptions flag all
user@host# set traceoptions level all
```

12. Go to the following hierarchy level and verify the configuration:

```
[edit security idp]
user@host# show
idp-policy test1 {
    rulebase-ips {
        rule r1 {
            match {
                application default;
                attacks {
                    predefined-attacks FTP:USER:ROOT;
                    predefined-attack-groups "Recommended Attacks";
                }
            }
            then {
                action {
                    no-action;
                }
                notification {
                    log-attacks {
                        alert;
                    }
                }
            }
        }
    }
}
```



```

}
traceoptions {
  file idp-demo.log;
  flag all;
  level all;
}

```

13. Go to the following hierarchy level:

```

[edit services]
user@host# edit aacl

```

14. In the hierarchy level, configure the AACL rules.

```

[edit services aacl]
user@host# set rule rule-name match-direction input-output term term from
  application-group-any
user@host# set rule rule-name match-direction input-output term term then count
  application accept

```

In this example, the AACL rule is **app-aware** and the term is **t1**.

```

[edit services aacl]
user@host# set rule app-aware match-direction input-output term t1 from
  application-group-any
user@host# set rule app-aware match-direction input-output term t1 then count
  application accept

```

15. Go to the following hierarchy level and verify the configuration:

```

[edit services aacl]
user@host# show
rule app-aware {
  match-direction input-output;
  term t1 {
    from {
      application-group-any;
    }
    then {
      count application;
      accept;
    }
  }
}

```

16. Go to the following hierarchy level:

```

[edit services]
user@host# edit service-set App-Aware-Set

```

17. Configure the APPID profile.

```

[edit services service-set App-Aware-Set]
user@host# set application-identification-profile application-identification-profile

```

In this example, the APPID profile is **dummy-profile**.

```

[edit services service-set App-Aware-Set]
user@host# set application-identification-profile dummy-profile

```

18. Configure the IDP profile.

```

[edit services service-set App-Aware-Set]

```

```
user@host# set idp-profile idp-profile
```

In this example, the IDP profile is **test1**.

```
[edit services service-set App-Aware-Set]
```

```
user@host# set idp-profile test1
```

19. Configure the policy decision statistics profile.

```
[edit services service-set App-Aware-Set]
```

```
user@host# set policy-decision-statistics-profile profile-name
```

In this example, the policy decision statistics profile is **lpdf-stats**.

```
[edit services service-set App-Aware-Set]
```

```
user@host# set policy-decision-statistics-profile lpdf-stats
```

20. Configure the AACL rules.

```
[edit services service-set App-Aware-Set]
```

```
user@host# set aacl-rules rule-name
```

In this example, the AACL rule name is **app-aware**.

```
[edit services service-set App-Aware-Set]
```

```
user@host# set aacl-rules app-aware
```

21. Configure two stateful firewall rules.

```
[edit services service-set App-Aware-Set]
```

```
user@host# set stateful-firewall-rules rule-name
```

```
user@host# set stateful-firewall-rules rule-name
```

In this example, the first rule is **r1** and the second rule is **r2**.

```
[edit services service-set App-Aware-Set]
```

```
user@host# set stateful-firewall-rules r1
```

```
user@host# set stateful-firewall-rules r2
```

22. In the hierarchy level, configure the service set to bypass traffic on service PIC failure.

```
[edit services service-set App-Aware-Set]
```

```
user@host# set service-set-options bypass-traffic-on-pic-failure
```

23. Configure interface-specific service set options.

```
[edit services service-set App-Aware-Set]
```

```
user@host# set interface-service service-interface service-interface
```

In this example, the services interface is **ms-0/1/0**.

```
[edit services service-set App-Aware-Set]
```

```
user@host# set interface-service service-interface ms-0/1/0
```

24. Go to the following hierarchy level and verify the configuration:

```
[edit services service-set App-Aware-Set]
```

```
user@host# show
```

```
application-identification-profile dummy-profile;
```

```
idp-profile test1;
```

```
policy-decision-statistics-profile {
```

```
  lpdf-stats;
```

```
}
```

```
aacl-rules app-aware;
```

```

stateful-firewall-rules r1;
stateful-firewall-rules r2;
service-set-options {
    bypass-traffic-on-pic-failure;
}
interface-service {
    service-interface ms-0/1/0;
}

```

25. Go to the following hierarchy level:

```

[edit services]
user@host# edit service-set NAT-SFW-SET

```

26. In the hierarchy level, configure optional notification parameters for the services interface. Note that it is required only for debugging.

```

[edit services service-set NAT-SFW-SET]
user@host# set syslog host host-name services any

```

In this example, the host to notify is **local**.

```

[edit services service-set NAT-SFW-SET]
user@host# set services-options syslog host local services any

```

27. Configure two stateful firewall rules.

```

[edit services service-set NAT-SFW-SET]
user@host# set stateful-firewall-rules rule-name
user@host# set stateful-firewall-rules rule-name

```

In this example, the first rule is **r1** and the second rule is **r2**.

```

[edit services service-set NAT-SFW-SET]
user@host# set stateful-firewall-rules r1
user@host# set stateful-firewall-rules r2

```

28. Configure NAT rules.

```

[edit services service-set NAT-SFW-SET]
user@host# set nat-rules rule-name

```

In this example, the NAT rule is **SET-MSR-ADDR**.

```

[edit services service-set NAT-SFW-SET]
user@host# set nat-rules SET-MSR-ADDR

```

29. Configure interface-specific service set options.

```

[edit services service-set NAT-SFW-SET]
user@host# set interface-service service-interface

```

In this example, the services interface is **sp-3/1/0**.

```

[edit services service-set NAT-SFW-SET]
user@host# set interface-service service-interface sp-3/1/0

```

30. Go to the following hierarchy level and verify the configuration:

```

[edit services service-set NAT-SFW-SET]
user@host# show
syslog {
    host local {
        services any;
    }
}

```

```
    }  
  }  
  stateful-firewall-rules r1;  
  stateful-firewall-rules r2;  
  interface-service {  
    service-interface sp-3/1/0;  
  }
```

31. Go to the following hierarchy level:

```
user@host# edit interfaces
```

32. In the hierarchy level, configure the interface.

```
[edit interfaces]  
user@host# set interface
```

In this example, the interface is **ge-1/2/2.1**.

```
[edit interfaces]  
user@host# set ge-1/2/2.1
```

33. Go to the following hierarchy level:

```
[edit interfaces]  
user@host# edit ge-1/2/2.1
```

34. In the hierarchy level, configure the service set for received packets.

```
[edit interfaces ge-1/2/2 unit 1]  
user@host# set family inet service input service-set service-set-name
```

In this example, the input service set is **App-Aware-Set**.

```
[edit interfaces ge-1/2/2 unit 1]  
user@host# set family inet service input service-set App-Aware-Set
```

35. Configure the service set for transmitted packets.

```
[edit interfaces ge-1/2/2 unit 1]  
user@host# set family inet service output service-set service-set-name
```

In this example, the output service set is **App-Aware-Set**.

```
[edit interfaces ge-1/2/2 unit 1]  
user@host# set family inet service output service-set App-Aware-Set
```

36. Go to the following hierarchy level:

```
[edit interfaces ge-1/2/2 unit 1]  
user@host# edit family inet
```

37. In the hierarchy level, configure the interface address.

```
[edit interfaces ge-1/2/2 unit 1 family inet]  
user@host# set address source
```

In this example, the interface address is **10.10.9.10/30**.

```
[edit interfaces]  
user@host# set address 10.10.9.10/30
```

38. Go to the following hierarchy level and verify the configuration:

```
[edit interfaces ge-1/2/2 unit 1]
user@host# show
family inet {
  service {
    input {
      service-set App-Aware-Set;
    }
    output {
      service-set App-Aware-Set;
    }
  }
  address 10.10.9.10/30;
}
```

39. Go to the following hierarchy level:

```
user@host# edit interfaces
```

40. In the hierarchy level, configure the interface.

```
[edit interfaces]
user@host# set interface
```

In this example, the interface is **ge-1/1/0.48**.

```
[edit interfaces]
user@host# set ge-1/1/0.48
```

41. Go to the following hierarchy level:

```
[edit interfaces]
user@host# edit ge-1/1/0.48
```

42. In the hierarchy level, configure the service set for received packets.

```
[edit interfaces ge-1/1/0 unit 48]
user@host# set family inet service input service-set service-set-name
```

In this example, the service set is **NAT-SFW-SET**.

```
[edit interfaces ge-1/1/0 unit 48]
user@host# set family inet service input service-set NAT-SFW-SET
```

43. Configure the service set for transmitted packets.

```
[edit interfaces ge-1/1/0 unit 48]
user@host# set family inet service output service-set service-set-name
```

In this example, the service set is **NAT-SFW-SET**.

```
[edit interfaces ge-1/1/0 unit 48]
user@host# set family inet service output service-set NAT-SFW-SET
```

44. Go to the following hierarchy level:

```
[edit interfaces ge-1/1/0 unit 48]
user@host# edit family inet
```

45. Configure the interface address.

```
[edit interfaces ge-1/1/0 unit 48 family inet]
user@host# set address source
```

In this example, the interface address is **10.48.0.1/31**.

```
[edit interfaces ge-1/1/0 unit 48 family inet]
user@host# set address 10.48.0.1/31
```

46. Go to the following hierarchy level and verify the configuration:

```
[edit interfaces ge-1/1/0 unit 48]
user@host# show
family inet {
  service {
    input {
      service-set NAT-SFW-SET;
    }
    output {
      service-set NAT-SFW-SET;
    }
  }
  address 10.48.0.1/31;
}
```

47. Go to the following hierarchy level:

```
user@host# edit interfaces
```

48. In the hierarchy level, configure the interface.

```
[edit interfaces]
set interface
```

In this example, the interface is **ms-0/1/0.0**.

```
[edit interfaces]
user@host# set ms-0/1/0.0
```

49. Go to the following hierarchy level:

```
[edit interfaces]
user@host# edit ms-0/1/0.0
```

50. In the hierarchy level, configure the protocol family.

```
[edit interfaces ms-0/1/0 unit 0]
user@host# set family inet
```

51. Go to the following hierarchy level and verify the configuration:

```
[edit interfaces ms-0/1/0]
user@host# show
unit 0 {
  family inet;
}
```

52. Go to the following hierarchy level:

```
user@host# edit interfaces
```

53. In the hierarchy level, configure the interface.

```
[edit interfaces]
set interface
```

In this example, the interface is **sp-3/1/0.0**.

```
[edit interfaces]
```

```
user@host# set sp-3/1/0.0
```

54. Go to the following hierarchy level:

```
[edit interfaces]
user@host# edit sp-3/1/0
```

55. In the hierarchy level, configure optional notification parameters for the services interface. Note that it is required only for debugging.

```
[edit interfaces sp-3/1/0]
user@host# set services-options syslog host host-name services any
```

In this example, the host to notify is **local**.

```
[edit interfaces sp-3/1/0]
user@host# set services-options syslog host local services any
```

56. Go to the following hierarchy level:

```
[edit interfaces]
user@host# edit sp-3/1/0.0
```

57. In the hierarchy level, configure the protocol family.

```
[edit interfaces sp-3/1/0 unit 0]
user@host# set family inet
```

58. Go to the following hierarchy level and verify the configuration:

```
[edit interfaces sp-3/1/0]
user@host# show
services-options {
  syslog {
    host local {
      services any;
    }
  }
}
unit 0 {
  family inet;
}
```

59. Go to the following hierarchy level:

```
[edit chassis]
```

60. In the hierarchy level, configure the redundancy settings.

```
[edit chassis]
user@host# set no-service-pic-restart-on-failover
user@host# set redundancy graceful-switchover
```

61. Configure the FPC and PIC.

```
[edit chassis]
user@host# edit fpc slot pic slot
```

In this example, the FPC is in slot 0 and the PIC is in slot 1.

```
[edit chassis]
user@host# edit fpc 0 pic 1
```

62. Configure the number of cores dedicated to run control functionality.

```
[edit chassis fpc slot pic slot]
user@host# set adaptive-services service-package extension-provider control-cores
control-cores
```

In this example, the number of control cores is 1.

```
[edit chassis fpc 1 pic 0]
user@host# set adaptive-services service-package extension-provider control-cores
1
```

63. Configure the number of processing cores dedicated to data.

```
[edit chassis fpc slot pic slot]
user@host# set adaptive-services service-package extension-provider data-cores
data-cores
```

In this example, the number of data cores is 7.

```
[edit chassis fpc 1 pic 0]
user@host# set adaptive-services service-package extension-provider data-cores 7
```

64. Configure the size of the object cache in megabytes. Only values in increments of 128 MB are allowed and the maximum value of object cache can be 1280 MB. On MS-100, the value is 512 MB.

```
[edit chassis fpc slot pic slot]
user@host# set adaptive-services service-package extension-provider
object-cache-size object-cache-size
```

In this example, the size of the object cache is 1280 MB.

```
[edit chassis fpc 1 pic 0]
user@host# set adaptive-services service-package extension-provider
object-cache-size 1280
```

65. Configure the size of the policy database in megabytes.

```
[edit chassis fpc slot pic slot]
user@host# set adaptive-services service-package extension-provider policy-db-size
policy-db-size
```

In this example, the size of the policy database is 64 MB.

```
[edit chassis fpc 1 pic 0]
user@host# set adaptive-services service-package extension-provider policy-db-size
64
```

66. Configure the packages.

```
[edit chassis fpc slot pic slot]
user@host# set adaptive-services service-package extension-provider package
package
```

In this example, the first package is **jservices-appid**, the second package is **jservices-aacl**, the third package is **jservices-llpdf**, the fourth package is **jservices-idp**, and the fifth package is **jservices-sfw**. **jservices-sfw** is available only in Junos OS Release 10.1 and later.

```
[edit chassis fpc 1 pic 0]
user@host# set adaptive-services service-package extension-provider package
jservices-appid
```



```

user@host# set adaptive-services service-package extension-provider package
jservices-aac1
user@host# set adaptive-services service-package extension-provider package
jservices-llpdf
user@host# set adaptive-services service-package extension-provider package
jservices-idp
user@host# set adaptive-services service-package extension-provider package
jservices-sfw

```

67. Configure the IP network services.

```

[edit chassis]
user@host# set network-services ip

```

68. Go to the following hierarchy level and verify the configuration:

```

[edit chassis]
user@host# show chassis
no-service-pic-restart-on-failover;
filter-memory-enhanced;
redundancy {
    graceful-switchover;
}
fpc 0 {
    pic 1 {
        adaptive-services {
            service-package {
                extension-provider {
                    control-cores 1;
                    data-cores 7;
                    object-cache-size 1280;
                    policy-db-size 64;
                    package jservices-appid;
                    package jservices-aac1;
                    package jservices-llpdf;
                    package jservices-idp;
                    package jservices-sfw;
                }
            }
        }
    }
}
network-services ip;

```

Example: Virtual Routing and Forwarding (VRF) and Service Configuration

The following example combines virtual routing and forwarding (VRF) and services configuration:

```

[edit policy-options]
policy-statement test-policy {
    term t1 {
        then reject;
    }
}
[edit routing-instances]
test {
    interface ge-0/2/0.0;
    interface sp-1/3/0.20;
}

```

```
instance-type vrf;
route-distinguisher 10.58.255.1:37;
vrf-import test-policy;
vrf-export test-policy;
routing-options {
    static {
        route 0.0.0.0/0 next-table inet.0;
    }
}
[edit interfaces]
ge-0/2/0 {
    unit 0 {
        family inet {
            service {
                input service-set nat-me;
                output service-set nat-me;
            }
        }
    }
}
sp-1/3/0 {
    unit 0 {
        family inet;
    }
    unit 20 {
        family inet;
        service-domain inside;
    }
    unit 21 {
        family inet;
        service-domain outside;
    }
}
[edit services]
stateful-firewall {
    rule allow-any-input {
        match-direction input;
        term t1 {
            then accept;
        }
    }
}
nat {
    pool hide-pool {
        address 10.58.16.100;
        port automatic;
    }
    rule hide-all-input {
        match-direction input;
        term t1 {
            then {
                translated {
                    source-pool hide-pool;
                    translation-type source napt-44;
                }
            }
        }
    }
}
```

```
    }  
  }  
}  
service-set nat-me {  
  stateful-firewall-rules allow-any-input;  
  nat-rules hide-all-input;  
  interface-service {  
    service-interface sp-1/3/0.20;  
  }  
}
```


CHAPTER 28

IDS Configuration Overview

- [Configuring IDS Rules on page 355](#)
- [Configuring IDS Rule Sets on page 363](#)
- [Examples: Configuring IDS Rules on page 364](#)

Configuring IDS Rules

IDS rules identify traffic for which you want the router software to count events. Because IDS is based on stateful firewall properties, you must configure at least one stateful firewall rule and include it in the service set with the IDS rules; for more information, see [“Configuring Stateful Firewall Rules” on page 331](#).

To configure an IDS rule, include the **rule** *rule-name* statement at the **[edit services ids]** hierarchy level:

```
[edit services ids]
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application-sets set-name;
      applications [ application-names ];
      destination-address (address | any-unicast) <except>;
      destination-address-range low minimum-value high maximum-value <except>;
      destination-prefix-list list-name <except>;
      source-address (address | any-unicast) <except>;
      source-address-range low minimum-value high maximum-value <except>;
      source-prefix-list list-name <except>;
    }
    then {
      aggregation {
        destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
        source-prefix prefix-value | source-prefix-ipv6 prefix-value;
      }
      (force-entry | ignore-entry);
      logging {
        syslog;
        threshold rate;
      }
      session-limit {
        by-destination {
```

```

        hold-time seconds;
        maximum number;
        packets number;
        rate number;
    }
    by-pair {
        hold-time seconds;
        maximum number;
        packets number;
        rate number;
    }
    by-source {
        hold-time seconds;
        maximum number;
        packets number;
        rate number;
    }
}
syn-cookie {
    mss value;
    threshold rate;
}
}
}

```

Each IDS rule consists of a set of terms, similar to a filter configured at the **[edit firewall]** hierarchy level. A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded.
- **then** statement—Specifies the actions and action modifiers to be performed by the router software.

The following sections describe IDS rule content in more detail:

- [Configuring Match Direction for IDS Rules on page 356](#)
- [Configuring Match Conditions in IDS Rules on page 357](#)
- [Configuring Actions in IDS Rules on page 358](#)

Configuring Match Direction for IDS Rules

Each rule must include a **match-direction** statement that specifies whether the match is applied on the input or output side of the interface. To configure where the match is applied, include the **match-direction (input | input-output | output)** statement at the **[edit services ids rule *rule-name*]** hierarchy level:

```

[edit services ids rule rule-name]
  match-direction (input | output | input-output);

```

If you configure **match-direction input-output**, bidirectional rule creation is allowed.

The match direction is used with respect to the traffic flow through the AS or Multiservices PIC. When a packet is sent to the PIC, direction information is carried along with it.

With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.

With a next-hop service set, packet direction is determined by the interface used to route the packet to the AS or Multiservices PIC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the PIC, the packet direction is output. For more information on inside and outside interfaces, see [“Configuring Service Sets to be Applied to Services Interfaces” on page 31](#).

On the AS or Multiservices PIC, a flow lookup is performed. If no flow is found, rule processing is performed. All rules in the service set are considered. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that match the packet direction are considered.

Configuring Match Conditions in IDS Rules

To configure IDS match conditions, include the **from** statement at the **[edit services ids rule rule-name term term-name]** hierarchy level:

```
[edit services ids rule rule-name term term-name]
from {
  application-sets set-name;
  applications [ application-names ];
  destination-address (address | any-unicast) <except>;
  destination-address-range low minimum-value high maximum-value <except>;
  destination-prefix-list list-name <except>;
  source-address (address | any-unicast) <except>;
  source-address-range low minimum-value high maximum-value <except>;
  source-prefix-list list-name <except>;
}
```

If you omit the **from** statement, the software accepts all events and places them in the IDS cache for processing.

The source address and destination address can be either IPv4 or IPv6. You can use the destination address, a range of destination addresses, a source address, or a range of source addresses as a match condition, in the same way that you would configure a firewall filter; for more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Alternatively, you can specify a list of source or destination prefixes by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level and then including either the **destination-prefix-list** or **source-prefix-list** statement in the IDS rule. For an example, see [“Examples: Configuring Stateful Firewall Rules” on page 335](#).

You can also include application protocol definitions that you have configured at the **[edit applications]** hierarchy level; for more information, see [“Configuring Application Protocol Properties” on page 303](#).

- To apply one or more specific application protocol definitions, include the **applications** statement at the **[edit services ids rule rule-name term term-name from]** hierarchy level.

- To apply one or more sets of application protocol definitions that you have defined, include the **application-sets** statement at the **[edit services ids rule *rule-name* term *term-name* from]** hierarchy level.



NOTE: If you include one of the statements that specifies application protocols, the router derives port and protocol information from the corresponding configuration at the **[edit applications]** hierarchy level; you cannot specify these properties as match conditions.

If a match occurs on an application, the application protocol is displayed separately in the **show services ids** command output. For more information, see the [CLI Explorer](#).

Configuring Actions in IDS Rules

To configure IDS actions, include the **then** statement at the **[edit services ids rule *rule-name* term *term-name*]** hierarchy level:

```
[edit services ids rule rule-name term term-name]
then {
  aggregation {
    destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
    source-prefix prefix-value | source-prefix-ipv6 prefix-value;
  }
  (force-entry | ignore-entry);
  logging {
    syslog;
    threshold rate;
  }
  session-limit {
    by-destination {
      hold-time seconds;
      maximum number;
      packets number;
      rate number;
    }
    by-pair {
      hold-time seconds;
      maximum number;
      packets number;
      rate number;
    }
    by-source {
      hold-time seconds;
      maximum number;
      packets number;
      rate number;
    }
  }
  syn-cookie {
    mss value;
    threshold rate;
  }
}
```


You can configure the following possible actions:

- **aggregation**—The router aggregates traffic labeled with the specified source or destination prefixes before passing the events to IDS processing. This is helpful if you want to examine all the traffic connected with a particular source or destination host. To collect traffic with some other marker, such as a particular application or port, configure that value in the match conditions.

To configure aggregation prefixes, include the **aggregation** statement at the **[edit services ids rule rule-name term term-name then]** hierarchy level and specify values for **source-prefix**, **destination-prefix**, **source-prefix-ipv6**, or **destination-prefix-ipv6**:

```
[edit services ids rule rule-name term term-name then]
  aggregation {
    destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
    source-prefix prefix-value | source-prefix-ipv6 prefix-value;
  }
```

The value of **source-prefix** and **destination-prefix** must be an integer between 1 and 32. The value of **source-prefix-ipv6** and **destination-prefix-ipv6** must be an integer between 1 and 128.

- **(force-entry | ignore-entry)**—**force-entry** provides a permanent spot in IDS caches for subsequent events after one event is registered. By default, the IDS software does not record information about “good” packets that do not exhibit suspicious behavior. You can use the **force-entry** statement to record all traffic from a suspect host, even traffic that would not otherwise be counted.

ignore-entry ensures that all IDS events are ignored. You can use this statement to disregard all traffic from a host you trust, including any temporary anomalies that IDS would otherwise count as events.

To configure an entry behavior different from the default, include the **force-entry** or **ignore-entry** statement at the **[edit services ids rule rule-name term term-name then]** hierarchy level:

```
[edit services ids rule rule-name term term-name then]
  (force-entry | ignore-entry);
```

- **logging**—The event is logged in the system log file.

To configure logging, include the **logging** statement at the **[edit services ids rule rule-name term term-name then]** hierarchy level:

```
[edit services ids rule rule-name term term-name then]
  logging {
    syslog;
    threshold rate;
  }
```

You can optionally include a threshold rate to trigger the generation of system log messages. The threshold rate is specified in events per second. IDS logs are generated once every 60 seconds for each anomaly that is reported. The logs are generated as long as the events continue.

- **session-limit**—The router limits open sessions when the specified threshold is reached.

To configure a threshold, include the **session-limit** statement at the **[edit services ids rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ids rule rule-name term term-name then]
session-limit {
  by-destination {
    hold-time seconds;
    maximum number;
    packets number;
    rate number;
  }
  by-pair {
    hold-time seconds;
    maximum number;
    packets number;
    rate number;
  }
  by-source {
    hold-time seconds;
    maximum number;
    packets number;
    rate number;
  }
}
```

You configure the thresholds for flow limitation based on traffic direction:

- To limit the number of outgoing sessions from one internal host or subnet, configure the **by-source** statement.
- To limit the number of sessions between a pair of IP addresses, subnets, or applications, configure the **by-pair** statement.
- To limit the number of incoming sessions to one external public IP address or subnet, configure the **by-destination** statement.

For each direction, you can configure the following threshold values:

- **hold-time *seconds***—When the **rate** or **packets** measurement reaches the threshold value, stop all new flows for the specified number of seconds. Once **hold-time** is in effect, the traffic is blocked for the specified time even if the rate subsides below the specified limit. By default, **hold-time** has a value of 0; the range is 0 through 60 seconds.
- **maximum *number***—Maximum number of open sessions per IP address or subnet per application. The range is 1 through 32,767.
- **packets *number***—Maximum number of packets per second (pps) per IP address or subnet per application. The range is 4 through 2,147,483,647.
- **rate *number***—Maximum number of sessions per second per IP address or subnet per application. The range is 4 through 32,767.

If you include more than one source address in the match conditions configured at the **[edit services ids rule *rule-name* term *term-name* from]** hierarchy level, limits are applied for each source address independently. For example, the following

configuration allows 20 connections from each source address (10.1.1.1 and 10.1.1.2), not 20 connections total. The same logic applies to the **applications** and **destination-address** match conditions.

```
[edit services ids rule rule-name term term-name]
  from {
    source-address 10.1.1.1;
    source-address 10.1.1.2;
  }
  then {
    session-limit by-source {
      maximum 20;
    }
  }
}
```



NOTE: IDS limits are applied to packets that are accepted by stateful firewall rules. They are not applied to packets discarded or rejected by stateful firewall rules. For example, if the stateful firewall accepts 75 percent of the incoming traffic and the remaining 25 percent is rejected or discarded, the IDS limit applies only to 75 percent of the traffic.

- **syn-cookie**—The router activates SYN-cookie defensive mechanisms.

To configure SYN-cookie values, include the **syn-cookie** statement at the **[edit services ids rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ids rule rule-name term term-name then]
  syn-cookie {
    mss value;
    threshold rate;
  }
```

If you enable SYN-cookie defenses, you must include both a threshold rate to trigger SYN-cookie activity and a Transmission Control Protocol (TCP) maximum segment size (MSS) value for TCP delayed binding. The threshold rate is specified in SYN attacks per second. By default, the TCP MSS value is 1500; the range is from 128 through 8192.

Handling of SYN Flood Attacks and SYN Cookie Protection

The main purpose of a SYN flood attack is to consume all new network connections at a site and thereby prevent authorized and legitimate users from being able to connect to network resources. The SYN (synchronize sequence number) packet is the first request to connect sent to a system. The SYN packet contains an ID to which the receiver is required to respond. If the packet contains an illegal ID, the receiving system does receive a connection acknowledgment when it responds to the intended connection initiator. Eventually, this half-open connection times out and the incoming channel on the receiver becomes available again to normally handle another request. A SYN flood attack sends so many such requests that all incoming connections are continuously tied up waiting for acknowledgments that are never received. This condition causes the server to be unavailable to legal users (except in cases where a user session is established when it is exactly at the moment when one of the tied-up connections times out). A SYN flood attack is a connectionless attack. It does not require a real source IP addresses and, because it uses legitimate destination IP or port addresses, is practically impossible to distinguish from legitimate packets. Therefore, it is very difficult to prevent this type of

attack by using only filters or stateful firewall rules. Basically, there are only three methods to protect from this type of attack:

- Intercept (delayed binding)—The firewall intercepts incoming TCP synchronization requests and establishes a connection with the client on the server's behalf, and with the server on the client's behalf. If both connections are successful, the firewall transparently merges the two connections. The firewall usually has aggressive timeouts to prevent its own resources from being consumed by a SYN attack. This is the most intensive solution in terms of processing and memory requirements.
- Watch (SYN defense)—The firewall passively watches half-open connections and actively closes connections on the server after a configurable length of time.
- SYN cookie—SYN cookies are particular choices for the initial TCP sequence number chosen by the TCP server. A host requesting a connection must answer with the cookie to connect to an open TCP socket while a SYN-flood has been detected as in progress by the IDS.

Juniper Networks routers support the combination of stateful firewall and IDS mechanisms to support the SYN cookie and watch (SYN defense) methods. The key to the SYN flood attack is the filling of the SYN queue of the victim or the attacked network element. The SYN cookie defense method enables the victim to continue accepting connection requests when the SYN queue is full or, in the case of the firewall or IDS applications, when a certain threshold has been reached. After the threshold is reached, a cryptographic cookie (a 32-bit number) is created from information in the SYN segment and the SYN segment is dropped. The cookie is used as the initial sequence number in the SYN-ACK sent to the client. The cookie (plus one) is returned to the firewall or IDS application as the acknowledgment number in the ACK from a legitimate client. The returned cookie can be validated and the most important parts of the SYN segment can be reconstructed from the cookie, thereby allowing a connection to be established. Because the spoofed clients of the SYN flood never send ACKs, no resources are allocated for them in any state when SYN cookies are in use. It is preferred that you use SYN flood countermeasures only for hosts under attack. The anomaly table can be used for reliable attack recognition or they can be enabled within the stateful firewall. Such a type of configuration also helps prevent the depletion of system resources (especially the flow table) in case of attacks.

When combining multiple services, the general path is an important factor for consideration in the forward and reverse directions. This is especially true when NAT is deployed to determine whether the pre-NAT or post-NAT address must be used to match a rule. In the forward path from a LAN interface to a WAN interface, IDS and stateful firewall are performed first, then NAT, and finally IPsec. This sequence of processing of services denotes that the stateful firewall must match on a pre-NAT address, whereas the IPsec tunnel matches on the post-NAT address. In the return path, the IPsec packet is processed first, then NAT, and finally the stateful firewall. This order of processing still allows IPsec to match a public address and the stateful firewall to match on a private address. You must separately configure the firewall, NAT, and IDS services. The processing of packets becomes much more complicated when IPsec over GRE is implemented in the router with other services turned on. This behavior occurs because Junos OS treats GRE packets in a unique fashion after GRE encapsulation. After a packet is encapsulated in a GRE packet, it is marked with an input interface as the next-hop outgoing interface.

This method of marking causes GRE packets to be blocked if any input filters or input services are allowed that do not allow for this service.

Junos OS services support a limited set of IDS rules to help detect attacks such as port scanning and anomalies in traffic patterns. It also supports some attack prevention by limiting the number of flows, sessions, and rates. In addition, it protects against SYN attacks by implementing a SYN cookie mechanism. Because the intrusion detection and prevention (IDP) service does not support higher-layer application signatures, an effective approach against attacks is that protection against a SYN attack can be configured. The IDP solution is largely a monitoring tool and not an essential prevention tool. To prevent a SYN attack, the router will operate as a type of SYN “proxy” and utilizes cookie values. When this feature is turned on, the router responds to the initial SYN packet with a SYN-ACK packet that contains a unique cookie value in the sequence number field. If the initiator responds with the same cookie in the sequence field, the TCP flow is accepted; if the responder does not respond or if it responds with the wrong cookie, the flow is dropped. To trigger this defense, you must configure a SYN cookie threshold. To enable the SYN cookie defense, an IDS rule action must contain a threshold that indicates when the feature should be enabled and an MSS value to avoid having the router manage segmented fragments when acting as a SYN proxy:

```
[edit]
user@host# set services ids rule simple-ids term 1 then syn-cookie
```

- Related Documentation**
- [Configuring IDS Rule Sets on page 363](#)
 - [Examples: Configuring IDS Rules on page 364](#)

Configuring IDS Rule Sets

The **rule-set** statement defines a collection of IDS rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services ids]** hierarchy level with a **rule** statement for each rule:

```
[edit services ids]
rule-set rule-set-name {
  rule rule-name;
}
```

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

- Related Documentation**
- [Configuring IDS Rules on page 355](#)
 - [Examples: Configuring IDS Rules on page 364](#)

Examples: Configuring IDS Rules

The following configuration adds a permanent entry to the IDS anomaly table when it encounters a flow with the destination address 10.410.6.2:

```
[edit services ids]
rule simple_ids {
  term 1 {
    from {
      destination-address 10.410.6.2/32;
    }
    then {
      force-entry;
      logging {
        threshold 1;
        syslog;
      }
    }
  }
  term default {
    then {
      aggregation {
        source-prefix 24;
      }
    }
  }
  match-direction input;
}
```

The IDS configuration works in conjunction with the stateful firewall mechanism and relies heavily on the anomalies reported by the stateful firewall. The following configuration example shows this relationship:

```
[edit services ids]
rule simple_ids {
  term 1 {
    from {
      source-address 10.30.20.2/32;
      destination-address {
        10.30.10.2/32;
        10.30.1.2/32 except;
      }
      applications appl-ftp;
    }
    then {
      force-entry;
      logging {
        threshold 5;
        syslog;
      }
      syn-cookie {
        threshold 10;
      }
    }
  }
}
```

```

    }
    match-direction input;
}

[edit services stateful-firewall]
rule my-firewall-rule {
  match-direction input-output;
  term term1 {
    from {
      source-address 10.30.20.2/32;
      applications appl-ftp ;
      destination-address {
        10.30.10.2/32;
        10.30.1.2/32 except;
      }
    }
    then {
      accept;
      syslog;
    }
  }
}

```

The stateful firewall or NAT service is used to generate the input data for the IDS application. When you enable and configure an IDS service, you must also enable stateful firewall with at least one rule (accept or discard all traffic). When the system is under an attack, the stateful firewall sends the correct and complete list of attack events to the IDS system. In your network environment, you can ensure that the system is wholly protected against a whole range of attacks so that the IDS system reports all such attacks. You must exercise caution when you configure the system to be protected from all attacks and unauthenticated access scenarios so that the traffic bandwidth that the system handles is not burdened. It is also important to verify the correlation between the firewall syslog messages corresponding to the attacks and IDS tables. The IDS tables must have the same or slightly less number of anomalies or errors compared to the firewall-based syslog messages. You can use the appropriate show commands are used to display the IDS tables.

A default stateful firewall rule can be as simple as only allowing connection initiation from the inside interface to the outside interface and discarding all other packets. However, in a real-world network environment, rules are generally more complex, such as configuring only a certain tributary unit ports are allowed to be opened, using application layer gateways (ALGs) for complicated protocols, and using NAT for both outgoing connections and inside hosts such as HTTP servers. Therefore, it is necessary to also configure the system as needed to interwork with simple and complicated rules. For example, if a SYN attack is directed towards an inside address that is simply discarded, no anomalies need to be reported to the IDS system. But if the SYN attack is directed towards the real HTTP server, anomalies must be reported. The IDS system can mitigate SYN attacks by using the TCP SYN cookie defense capability. You can enable the SYN cookie protection methodology by setting a threshold for SYNs per second for a given host and also a maximum segment size (MSS). Because the IDS system uses the stateful firewall, a firewall rule must be defined in the service-set. If you do not configure the **from** statement in a stateful firewall (rule term match condition) at the **[edit services service-set**

service-set-name stateful-firewall-rules *rule-name* term *term-name*] hierarchy level, it signifies that all events are placed into the IDS cache.

The following example shows configuration of flow limits:

```
[edit services ids]
rule ids-all {
  match-direction input;
  term t1 {
    from {
      application-sets alg-set;
    }
    then {
      aggregation {
        destination-prefix 30; /* IDS action aggregation */
      }
      logging {
        threshold 10;
      }
      session-limit {
        by-destination {
          hold-time 0;
          maximum 10;
          packets 200;
          rate 100;
        }
        by-pair {
          hold-time 0;
          maximum 10;
          packets 200;
          rate 100;
        }
        by-source {
          hold-time 5;
          maximum 10;
          packets 200;
          rate 100;
        }
      }
    }
  }
}
```

- Related Documentation**
- [Configuring IDS Rules on page 355](#)
 - [Configuring IDS Rule Sets on page 363](#)

CHAPTER 29

Monitoring Junos Network Secure

- [Monitoring Stateful Firewall Conversations on page 367](#)
- [Monitoring CGN, Stateful Firewall, and Software Flows on page 367](#)
- [Monitoring Global Stateful Firewall Statistics on page 368](#)

Monitoring Stateful Firewall Conversations

Purpose Use the `show services stateful-firewall conversations` command to show conversations, or collections of related flows.

Action `user@host# show services stateful-firewall conversations`
Interface: sp-0/0/0, Service set: sset
Conversation: ALG protocol: tcp
Number of initiators: 1, Number of responders: 1
Flow State Dir Frm
count
TCP 10.0.0.1:1025 -> 128.0.0.1:80 Forward I 372755
NAT source 10.0.0.1:1025 -> 129.0.0.1:1024
Software 2001:0:0:1::1 -> 1001::1
TCP 128.0.0.1:80 -> 129.0.0.1:1024 Forward O 794083
NAT dest 129.0.0.1:1024 -> 10.0.0.1:1025
Software 2001:0:0:1::1 -> 1001::1

Monitoring CGN, Stateful Firewall, and Software Flows

Purpose Use the following commands to check the creation of the softwires, pre-NAT flows, and post-NAT flows. Output can be filtered using more specific fields such as AFTR or B4 address or both for DS-Lite, and software-concentrator or software-initiator or both for 6rd.

- `show services stateful-firewall flows`
- `show services software flows`

Action user@host# **show services stateful-firewall flows**

Interface: sp-0/1/0, Service set: dslite-svc-set2

Flow	State	Dir	Frm count
TCP 200.200.200.2:80 -> 44.44.44.1:1025	Forward	O	219942
NAT dest 44.44.44.1:1025 -> 20.20.1.4:1025			
Software 2001::2 -> 1001::1			
TCP 20.20.1.2:1025 -> 200.200.200.2:80	Forward	I	110244
NAT source 20.20.1.2:1025 -> 44.44.44.1:1024			
Software 2001::2 -> 1001::1			
TCP 200.200.200.2:80 -> 44.44.44.1:1024	Forward	O	219140
NAT dest 44.44.44.1:1024 -> 20.20.1.2:1025			
Software 2001::2 -> 1001::1			
DS-LITE 2001::2 -> 1001::1	Forward	I	988729
TCP 200.200.200.2:80 -> 44.44.44.1:1026	Forward	O	218906
NAT dest 44.44.44.1:1026 -> 20.20.1.3:1025			
Software 2001::2 -> 1001::1			
TCP 20.20.1.3:1025 -> 200.200.200.2:80	Forward	I	110303
NAT source 20.20.1.3:1025 -> 44.44.44.1:1026			
Software 2001::2 -> 1001::1			
TCP 20.20.1.4:1025 -> 200.200.200.2:80	Forward	I	110944
NAT source 20.20.1.4:1025 -> 44.44.44.1:1025			
Software 2001::2 -> 1001::1			

- Related Documentation**
- *NAT Objects MIB in SNMP MIBs and Traps Reference*
 - *Network Address Translation Resources—Monitoring MIB in SNMP MIBs and Traps Reference*
 - *Juniper Networks Enterprise-Specific NAT Traps on SRX Series Services Gateways in SNMP MIBs and Traps Reference*

Monitoring Global Stateful Firewall Statistics

Purpose Use the **show services stateful-firewall statistics** command to observe statistics for service sets containing software rules.

Action user@host# **show services stateful-firewall statistics**

Interface Service set Accept Discard Reject Errors

sp-0/0/0 dslite-svc-set2 118991296 0 0 0

sp-0/1/0 dslite-svc-set1 237615050 0 0 0

PART 7

Creating Secure Tunnels Using Junos VPN Site Secure

- [Junos VPN Site Secure Overview on page 371](#)
- [Junos VPN Site Secure Configuration Overview on page 383](#)
- [Enhancing Security with Static IPsec over VRF on page 447](#)
- [Dynamically Assigning Tunnels Using Junos VPN Site Secure on page 455](#)
- [Enabling IPsec for the Services SDK on page 507](#)

CHAPTER 30

Junos VPN Site Secure Overview

- [Understanding Junos VPN Site Secure on page 371](#)
- [Authentication Algorithms on page 374](#)
- [Encryption Algorithms on page 374](#)
- [IPsec Protocols on page 376](#)
- [Supported IPsec and IKE Standards on page 378](#)
- [IPsec Terms and Acronyms on page 379](#)

Understanding Junos VPN Site Secure

Junos VPN Site Secure is a suite of IPsec features supported on multiservices line cards (MS-DPC, MS-MPC, and MS-MIC), and was referred to as IPsec services in Junos releases earlier than 13.2. In Junos OS Release 13.2 and later, the term IPsec features is used exclusively to refer to the IPsec implementation on Adaptive Services and Encryption Services PICs. This topic provides you an overview of Junos VPN Site Secure, and has the following sections:

- [IPsec on page 371](#)
- [Security Associations on page 372](#)
- [IKE on page 372](#)
- [Comparison of IPsec on ES PICs and Junos VPN Site Secure on Multiservices Line Cards on page 372](#)

IPsec

The IPsec architecture provides a security suite for the IP version 4 (IPv4) and IP version 6 (IPv6) network layers. The suite provides such functionality as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. In addition to IPsec, the Junos OS also supports the Internet Key Exchange (IKE), which defines mechanisms for key generation and exchange, and manages security associations (SAs).

IPsec also defines a security association and key management framework that can be used with any network-layer protocol. The SA specifies what protection policy to apply to traffic between two IP-layer entities. IPsec provides secure tunnels between two peers.

Security Associations

To use IPsec security services, you create SAs between hosts. An SA is a simplex connection that enables two hosts to communicate with each other securely by means of IPsec. There are two types of SAs:

- Manual SAs require no negotiation; all values, including the keys, are static and specified in the configuration. Manual SAs statically define the security parameter index (SPI) values, algorithms, and keys to be used, and require matching configurations on both ends of the tunnel. Each peer must have the same configured options for communication to take place.
- Dynamic SAs require additional configuration. With dynamic SAs, you configure IKE first and then the SA. IKE creates dynamic security associations; it negotiates SAs for IPsec. The IKE configuration defines the algorithms and keys used to establish the secure IKE connection with the peer security gateway. This connection is then used to dynamically agree upon keys and other data used by the dynamic IPsec SA. The IKE SA is negotiated first and then used to protect the negotiations that determine the dynamic IPsec SAs.

IKE

IKE is a key management protocol that creates dynamic SAs; it negotiates SAs for IPsec. An IKE configuration defines the algorithms and keys used to establish a secure connection with a peer security gateway.

IKE performs the following tasks:

- Negotiates and manages IKE and IPsec parameters.
- Authenticates secure key exchange.
- Provides mutual peer authentication by means of shared secrets (not passwords) and public keys.
- Provides identity protection (in main mode).

Two versions of the IKE protocol (IKEv1 and IKEv2) are supported now. IKE negotiates security attributes and establishes shared secrets to form the bidirectional IKE SA. In IKE, inbound and outbound IPsec SAs are established and the IKE SA secures the exchanges. Starting with Junos OS Release 11.4, both IKEv1 and IKEv2 are supported by default on all M Series, MX Series, and T Series routers. IKE also generates keying material, provides Perfect Forward Secrecy, and exchanges identities.

Comparison of IPsec on ES PICs and Junos VPN Site Secure on Multiservices Line Cards

[Table 23 on page 373](#) compares the top-level configuration of IPsec features on the ES PIC interfaces, and IPsec on the Adaptive Services PICs and Junos VPN Site Secure on Multiservices Line Cards .

Table 23: Statement Equivalents for ES and AS Interfaces

ES PIC Configuration	AS and MultiServices Line Cards Configuration
[edit security ipsec] proposal {...}	[edit services ipsec-vpn ipsec] proposal {...}
[edit security ipsec] policy {...}	[edit services ipsec-vpn ipsec] policy {...}
[edit security ipsec] security-association sa-dynamic {...}	[edit services ipsec-vpn rule <i>rule-name</i>] term <i>term-name</i> match-conditions {...} then dynamic {...}
[edit security ipsec] security-association sa-manual {...}	[edit services ipsec-vpn rule <i>rule-name</i>] term <i>term-name</i> match-conditions {...} then manual {...}
[edit security ike] proposal {...}	[edit services ipsec-vpn ike] proposal {...}
[edit security ike] policy {...}	[edit services ipsec-vpn ike] policy {...}
Not available	[edit services ipsec-vpn] rule-set {...}
Not available	[edit services ipsec-vpn] service-set {...}
[edit interfaces <i>es-fpc/pic/port</i>] tunnel source <i>address</i>	[edit services ipsec-vpn service-set <i>set-name</i>] ipsec-vpn local-gateway <i>address</i>
[edit interfaces <i>es-fpc/pic/port</i>] tunnel destination <i>address</i>	[edit services ipsec-vpn rule <i>rule-name</i>] remote-gateway <i>address</i>



NOTE: Although many of the same statements and properties are valid on both platforms (MultiServices and ES), the configurations are not interchangeable. You must commit a complete configuration for the PIC type that is installed in your router.

Related Documentation

- [Authentication Algorithms on page 374](#)
- [Encryption Algorithms on page 374](#)
- [IPsec Protocols on page 376](#)
- [Service Sets for IPsec Tunnels on page 430](#)
- [Configuring Security Associations on page 385](#)
- [IPsec Hierarchy Level on page 1297](#)

Authentication Algorithms

Authentication is the process of verifying the identity of the sender. Authentication algorithms use a shared key to verify the authenticity of the IPsec devices. The Junos OS uses the following authentication algorithms:

- Message Digest 5 (MD5) uses a one-way hash function to convert a message of arbitrary length to a fixed-length message digest of 128 bits. Because of the conversion process, it is mathematically infeasible to calculate the original message by computing it backwards from the resulting message digest. Likewise, a change to a single character in the message will cause it to generate a very different message digest number.

To verify that the message has not been tampered with, the Junos OS compares the calculated message digest against a message digest that is decrypted with a shared key. The Junos OS uses the MD5 hashed message authentication code (HMAC) variant that provides an additional level of hashing. MD5 can be used with authentication header (AH), Encapsulating Security Payload (ESP), and Internet Key Exchange (IKE).

- Secure Hash Algorithm 1 (SHA-1) uses a stronger algorithm than MD5. SHA-1 takes a message of less than 264 bits in length and produces a 160-bit message digest. The large message digest ensures that the data has not been changed and that it originates from the correct source. The Junos OS uses the SHA-1 HMAC variant that provides an additional level of hashing. SHA-1 can be used with AH, ESP, and IKE.
- SHA-256, SHA-384, and SHA-512 (sometimes grouped under the name SHA-2) are variants of SHA-1 and use longer message digests. The Junos OS supports the SHA-256 version of SHA-2, which can process all versions of Advanced Encryption Standard (AES), Data Encryption Standard (DES), and Triple DES (3DES) encryption.

Related Documentation

- [Understanding Junos VPN Site Secure on page 371](#)
- [Encryption Algorithms on page 374](#)

Encryption Algorithms

Encryption encodes data into a secure format so that it cannot be deciphered by unauthorized users. Like authentication algorithms, a shared key is used with encryption algorithms to verify the authenticity of the IPsec devices. The Junos OS uses the following encryption algorithms:

- Data Encryption Standard cipher-block chaining (DES-CBC) is a symmetric secret-key block algorithm. DES uses a key size of 64 bits, where 8 bits are used for error detection and the remaining 56 bits provide encryption. DES performs a series of simple logical operations on the shared key, including permutations and substitutions. CBC takes the first block of 64 bits of output from DES, combines this block with the second block, feeds this back into the DES algorithm, and repeats this process for all subsequent blocks.
- Triple DES-CBC (3DES-CBC) is an encryption algorithm that is similar to DES-CBC, but provides a much stronger encryption result because it uses three keys for 168-bit

(3 x 56-bit) encryption. 3DES works by using the first key to encrypt the blocks, the second key to decrypt the blocks, and the third key to re-encrypt the blocks.

- Advanced Encryption Standard (AES) is a next-generation encryption method based on the Rijndael algorithm developed by Belgian cryptographers Dr. Joan Daemen and Dr. Vincent Rijmen. It uses a 128-bit block and three different key sizes (128, 192, and 256 bits). Depending on the key size, the algorithm performs a series of computations (10, 12, or 14 rounds) that include byte substitution, column mixing, row shifting, and key addition. The use of AES in conjunction with IPsec is defined in RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*.

**Related
Documentation**

- [Understanding Junos VPN Site Secure on page 371](#)
- [Configuring IKE Proposals on page 405](#)
- [Configuring IPsec Proposals on page 415](#)
- [encryption on page 1357](#)

IPsec Protocols

IPsec protocols determine the type of authentication and encryption applied to packets that are secured by the router. The Junos OS supports the following IPsec protocols:

- **AH**—Defined in RFC 2402, AH provides connectionless integrity and data origin authentication for IPv4 and IPv6 packets. It also provides protection against replays. AH authenticates as much of the IP header as possible, as well as the upper-level protocol data. However, some IP header fields might change in transit. Because the value of these fields might not be predictable by the sender, they cannot be protected by AH. In an IP header, AH can be identified with a value of **51** in the **Protocol** field of an IPv4 packet and the **Next Header** field of an IPv6 packet. An example of the IPsec protection offered by AH is shown in [Figure 19 on page 376](#).



NOTE: AH is not supported on the T Series, M120, and M320 routers.

Figure 19: AH Protocol

Header format

Byte 0	Byte 1	Byte 2	Byte 3
Next header	Payload length	Reserved	
Security Parameters Index (SPI)			
Sequence number			
Authentication data (variable)			

Original IPv4 packet before AH is applied

Original IP header	TCP header	Data
--------------------	------------	------

IPv4 packet after AH transport mode is applied

Original IP header	AH header	TCP header	Data
Authenticating			


IPv4 packet after AH tunnel mode is applied

New IP header	AH header	Original IP header	TCP header	Data
Authenticating				

g015522

- ESP—Defined in RFC 2406, ESP can provide encryption and limited traffic flow confidentiality, or connectionless integrity, data origin authentication, and an anti-replay service. In an IP header, ESP can be identified a value of **50** in the **Protocol** field of an IPv4 packet and the **Next Header** field of an IPv6 packet. An example of the IPsec protection offered by ESP is shown in [Figure 20 on page 377](#).

Figure 20: ESP Protocol



- Bundle—When you compare AH with ESP, there are some benefits and shortcomings in both protocols. ESP provides a decent level of authentication and encryption, but does so only for part of the IP packet. Conversely, although AH does not provide encryption, it does provide authentication for the entire IP packet. Because of this, the Junos OS offers a third form of IPsec protocol called a protocol bundle. The bundle option offers a hybrid combination of AH authentication with ESP encryption.

- Related Documentation**
- [Understanding Junos VPN Site Secure on page 371](#)
 - [Configuring IPsec Proposals on page 415](#)
 - [Configuring Security Associations on page 385](#)
 - [protocol \(IPSec\) on page 1440](#)

Supported IPsec and IKE Standards

On routers equipped with one or more Adaptive Services PICs (both standalone and integrated versions) or Multiservices PICs or DPCs, the Canada and U.S. version of Junos OS substantially supports the following RFCs, which define standards for IP Security (IPsec) and Internet Key Exchange (IKE).

- RFC 2085, *HMAC-MD5 IP Authentication with Replay Prevention*
- RFC 2401, *Security Architecture for the Internet Protocol* (obsoleted by RFC 4301)
- RFC 2402, *IP Authentication Header* (obsoleted by RFC 4302)

This RFC is not supported on the ES PIC.

- RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*
- RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH* (obsoleted by RFC 4305)
- RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*
- RFC 2406, *IP Encapsulating Security Payload (ESP)* (obsoleted by RFC 4303 and RFC 4305)
- RFC 2407, *The Internet IP Security Domain of Interpretation for ISAKMP* (obsoleted by RFC 4306)
- RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)* (obsoleted by RFC 4306)
- RFC 2409, *The Internet Key Exchange (IKE)* (obsoleted by RFC 4306)
- RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*
- RFC 2451, *The ESP CBC-Mode Cipher Algorithms*
- RFC 2460, *Internet Protocol, Version 6 (IPv6)*
- RFC 3193, *Securing L2TP using IPsec*
- RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*
- RFC 3947, *Negotiation of NAT-Traversal in the IKE*
- RFC 3948, *UDP Encapsulation of IPsec ESP Packets*
- RFC 4301, *Security Architecture for the Internet Protocol*
- RFC 4302, *IP Authentication Header*

This RFC is not supported on the ES PIC.

- RFC 4303, *IP Encapsulating Security Payload (ESP)*
- RFC 4305, *Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)*
- RFC 4306, *Internet Key Exchange (IKEv2) Protocol*
- RFC 4307, *Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*

- RFC 4308, *Cryptographic Suites for IPsec*



NOTE: Only Suite VPN-A is supported in Junos OS.

- RFC 4835, *Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)*
- RFC 5996, *Internet Key Exchange Protocol Version 2 (IKEv2)*

Junos OS partially supports the following RFCs for IPsec and IKE:

- RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)*
- RFC 5114, *Additional Diffie-Hellman Groups for Use with IETF Standards*
- RFC 5903, *Elliptic Curve Groups modulo a Prime (ECP Groups) for IKE and IKEv2*

The following RFCs and Internet draft do not define standards, but provide information about IPsec, IKE, and related technologies. The IETF classifies them as “Informational.”

- RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*
- RFC 2412, *The OAKLEY Key Determination Protocol*
- RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*
- Internet draft draft-eastlake-sha2-02.txt, *US Secure Hash Algorithms (SHA and HMAC-SHA)* (expires July 2006)

Related Documentation

- [Services Interfaces Overview for Routing Devices](#)
- [MX Series Interface Module Reference](#)
- [Accessing Standards Documents on the Internet](#)

IPsec Terms and Acronyms

A

Adaptive Services PIC	A next-generation Physical Interface Card (PIC) that provides IPSec services and other services, such as Network Address Translation (NAT) and stateful firewall, on M Series and T Series platforms.
Advanced Encryption Standard (AES)	A next-generation encryption method that is based on the Rijndael algorithm and uses a 128-bit block, three different key sizes (128, 192, and 256 bits), and multiple rounds of processing to encrypt data.
authentication header (AH)	A component of the IPSec protocol used to verify that the contents of a packet have not changed (data integrity), and to validate the identity of the sender (data source authentication). For more information about AH, see RFC 2402.

C

- certificate authority (CA)** A trusted third-party organization that generates, enrolls, validates, and revokes digital certificates. The CA guarantees the identity of a user and issues public and private keys for message encryption and decryption.
- certificate revocation list (CRL)** A list of digital certificates that have been invalidated before their expiration date, including the reasons for their revocation and the names of the entities that have issued them. A CRL prevents usage of digital certificates and signatures that have been compromised.
- cipher block chaining (CBC)** A cryptographic method that encrypts blocks of ciphertext by using the encryption result of one block to encrypt the next block. Upon decryption, the validity of each block of ciphertext depends on the validity of all the preceding ciphertext blocks. For more information on how to use CBC with DES and ESP to provide confidentiality, see RFC 2405.

D

- Data Encryption Standard (DES)** An encryption algorithm that encrypts and decrypts packet data by processing the data with a single shared key. DES operates in increments of 64-bit blocks and provides 56-bit encryption.
- digital certificate** Electronic file that uses private and public key technology to verify the identity of a certificate creator and distribute keys to peers.

E

- Encapsulating Security Payload (ESP)** A component of the IPSec protocol used to encrypt data in an IPv4 or IPv6 packet, provide data integrity, and ensure data source authentication. For more information about ESP, see RFC 2406.
- ES PIC** A PIC that provides first-generation encryption services and software support for IPSec on M Series and T Series platforms.

H

- Hashed Message Authentication Code (HMAC)** A mechanism for message authentication using cryptographic hash functions. HMAC can be used with any iterative cryptographic hash function, such as MD5 or SHA-1, in combination with a secret shared key. For more information on HMAC, see RFC 2104.

I

- Internet Key Exchange (IKE)** Establishes shared security parameters for any hosts or routers using IPSec. IKE establishes the SAs for IPSec. For more information about IKE, see RFC 2407.

M

- Message Digest 5 (MD5)** An authentication algorithm that takes a data message of arbitrary length and produces a 128-bit message digest. For more information, see RFC 1321.

P

- Perfect Forward Secrecy (PFS)** Provides additional security by means of a Diffie-Hellman shared secret value. With PFS, if one key is compromised, previous and subsequent keys are secure because they are not derived from previous keys.

public key infrastructure (PKI) A trust hierarchy that enables users of a public network to securely and privately exchange data through the use of public and private cryptographic key pairs that are obtained and shared with peers through a trusted authority.

R

registration authority (RA) A trusted third-party organization that acts on behalf of a CA to guarantee the identity of a user.

Routing Engine A PCI-based architectural portion of a Junos OS-based router that handles the routing protocol process, the interface process, some of the chassis components, system management, and user access.

S

Secure Hash Algorithm 1 (SHA-1) An authentication algorithm that takes a data message of less than 264 bits in length and produces a 160-bit message digest. For more information on SHA-1, see RFC 3174.

Secure Hash Algorithm 2 (SHA-2) A successor to the SHA-1 authentication algorithm that includes a group of SHA-1 variants (SHA-224, SHA-256, SHA-384, and SHA-512). SHA-2 algorithms use larger hash sizes and are designed to work with enhanced encryption algorithms such as AES.

security association (SA) Specifications that must be agreed upon between two network devices before IKE or IPSec are allowed to function. SAs primarily specify protocol, authentication, and encryption options.

Security Association Database (SADB) A database where all SAs are stored, monitored, and processed by IPSec.

Security Parameter Index (SPI) An identifier that is used to uniquely identify an SA at a network host or router.

Security Policy Database (SPD) A database that works with the SADB to ensure maximum packet security. For inbound packets, IPSec checks the SPD to verify if the incoming packet matches the security configured for a particular policy. For outbound packets, IPSec checks the SPD to see if the packet needs to be secured.

Simple Certificate Enrollment Protocol (SCEP) A protocol that supports CA and registration authority (RA) public key distribution, certificate enrollment, certificate revocation, certificate queries, and certificate revocation list (CRL) queries.

T

Triple Data Encryption Standard (3DES) An enhanced DES algorithm that provides 168-bit encryption by processing data three times with three different keys.

CHAPTER 31

Junos VPN Site Secure Configuration Overview

- [Minimum Security Association Configurations on page 383](#)
- [Configuring Security Associations on page 385](#)
- [Example: Configuring Manual SAs on page 391](#)
- [Configuring IKE Proposals on page 405](#)
- [Configuring IKE Policies on page 409](#)
- [Configuring IPsec Proposals on page 415](#)
- [Configuring IPsec Policies on page 420](#)
- [Configuring IPsec Rules on page 422](#)
- [Configuring IPsec Rule Sets on page 429](#)
- [Service Sets for IPsec Tunnels on page 430](#)
- [Configuring IPsec Service Sets on page 430](#)
- [Tracing Junos VPN Site Secure Operations on page 436](#)
- [Multitask Example: Configuring IPsec Services on page 438](#)

Minimum Security Association Configurations

The following sections show the minimum configurations necessary to set up security associations (SAs) for IPsec services:

- [Minimum Manual SA Configuration on page 383](#)
- [Minimum Dynamic SA Configuration on page 384](#)

Minimum Manual SA Configuration

To define a manual SA configuration, you must include at least the following statements at the `[edit services ipsec-vpn rule rule-name term term-name then manual]` hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual]
  direction (inbound | outbound | bidirectional) {
    authentication {
      algorithm (hmac-md5-96 | hmac-sha1-96);
```

```

        key (ascii-text key | hexadecimal key);
    }
    encryption {
        algorithm algorithm;
        key (ascii-text key | hexadecimal key);
    }
    protocol (ah | esp | bundle);
    spi spi-value;
}

```

Minimum Dynamic SA Configuration

To define a dynamic SA configuration, you must include at least the following statements at the `[edit services ipsec-vpn]` hierarchy level:

```

[edit services ipsec-vpn]
ike {
    proposal proposal-name {
        authentication-algorithm (md5 | sha1 | sha-256);
        authentication-method pre-shared-keys;
        dh-group (group1 | group2 | group5 | group14);
        encryption-algorithm algorithm;
    }
    policy policy-name {
        proposals [ ike-proposal-names ];
        pre-shared-key (ascii-text key | hexadecimal key);
        version (1 | 2);
        mode (aggressive | main);
    }
}
ipsec {
    policy policy-name {
        proposals [ ipsec-proposal-names ];
    }
    proposal proposal-name {
        authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
        encryption-algorithm algorithm;
        protocol (ah | esp | bundle);
    }
}

```



NOTE:

- Starting with Junos OS Release 11.4, both IKEv1 and IKEv2 are supported by default on all M Series, MX Series, and T Series routers. The version statement at the `[edit services ipsec-vpn ike policy name]` hierarchy level allows you to configure the specific IKE version to be supported.
- The mode statement at the `[edit services ipsec-vpn ike policy name]` hierarchy level is required only if the version option is set to 1.

You must also include the `ipsec-policy` statement at the `[edit services ipsec-vpn rule rule-name term term-name then dynamic]` hierarchy level.

- Related Documentation**
- [Understanding Junos VPN Site Secure on page 371](#)
 - [Configuring Security Associations on page 385](#)
 - [Configuring IKE Proposals on page 405](#)
 - [Configuring IKE Policies on page 409](#)
 - [Configuring IPsec Proposals on page 415](#)
 - [Configuring IPsec Policies on page 420](#)

Configuring Security Associations

To use IPsec services, you create a security association (SA) between hosts. An SA is a simplex connection that enables two hosts to communicate with each other securely using IPsec.



NOTE: Both OSPFv2 and OSPFv3 support IPsec authentication. However, dynamic or tunnel mode IPsec SAs are not supported for OSPFv3. If you add SAs into OSPFv3 by including the `ipsec-sa` statement at the `[edit protocols ospf3 area area-number interface interface-name]` hierarchy level, your configuration commit fails. For more information about OSPF authentication and other OSPF properties, see the *Junos OS Routing Protocols Library for Routing Devices*.

You can configure two types of SAs:

- **Manual**—Requires no negotiation; all values, including the keys, are static and specified in the configuration. As a result, each peer must have the same configured options for communication to take place.
- **Dynamic**—Specifies proposals to be negotiated with the tunnel peer. The keys are generated as part of the negotiation and therefore do not need to be specified in the configuration. The dynamic SA includes one or more **proposal** statements that prioritizes a list of protocols and algorithms to be negotiated with the peer.

This section includes the following topics:

- [Configuring Manual Security Associations on page 385](#)
- [Configuring Dynamic Security Associations on page 390](#)
- [Clearing Security Associations on page 390](#)

Configuring Manual Security Associations

Manual SAs require no negotiation; all values, including the keys, are static and specified in the configuration. As a result, each peer must have the same configured options for communication to take place. Manual SAs are best suited for small, static networks where the distribution, maintenance, and tracking of keys are not difficult.

To configure a manual IPsec security association, include the following statements at the `[edit services ipsec-vpn rule rule-name term term-name then manual]` hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual]
direction (inbound | outbound | bidirectional) {
  authentication {
    algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
    key (ascii-text key | hexadecimal key);
  }
  auxiliary-spi auxiliary-spi-value;
  encryption {
    algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
    key (ascii-text key | hexadecimal key);
  }
  protocol (ah | esp | bundle);
  spi spi-value;
}
```

To configure manual SA statements, do the following:

- [Configuring the Direction for IPsec Processing on page 386](#)
- [Configuring the Protocol for a Manual IPsec SA on page 387](#)
- [Configuring the Security Parameter Index on page 387](#)
- [Configuring the Auxiliary Security Parameter Index on page 388](#)
- [Configuring Authentication for a Manual IPsec SA on page 388](#)
- [Configuring Encryption for a Manual IPsec SA on page 389](#)

Configuring the Direction for IPsec Processing

The **direction** statement specifies inbound or outbound IPsec processing. If you want to define different algorithms, keys, or security parameter index (SPI) values for each direction, you configure the **inbound** and **outbound** options. If you want the same attributes in both directions, use the **bidirectional** option.

To configure the direction of IPsec processing, include the **direction** statement at the `[edit services ipsec-vpn rule rule-name term term-name then manual]` hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual]
direction (inbound | outbound | bidirectional) {
  ...
}
```

The following two examples illustrate this:

- Example: Using Different Configuration for the Inbound and Outbound Directions

Define different algorithms, keys, and security parameter index values for each direction:

```
[edit services ipsec-vpn rule rule-name term term-name then manual]
direction bidirectional {
  protocol ah;
  spi 20001;
  authentication {
    algorithm hmac-md5-96;
```

```

        key ascii-text 123456789012abcd;
    }
}
direction outbound {
    protocol esp;
    spi 24576;
    encryption {
        algorithm 3des-cbc;
        key ascii-text 12345678901234567890abcd;
    }
}

```

- Example: Using the Same Configuration for the Inbound and Outbound Directions

Define one set of algorithms, keys, and security parameter index values that is valid in both directions:

```

[edit services ipsec-vpn rule rule-name term term-name then manual]
direction bidirectional {
    protocol ah;
    spi 20001;
    authentication {
        algorithm hmac-md5-96;
        key ascii-text 123456789012abcd;
    }
}

```

Configuring the Protocol for a Manual IPsec SA

IPsec uses two protocols to protect IP traffic: Encapsulating Security Payload (ESP) and authentication header (AH). The AH protocol is used for strong authentication. A third option, **bundle**, uses AH authentication and ESP encryption; it does not use ESP authentication because AH provides stronger authentication of IP packets.

To configure the IPsec protocol, include the **protocol** statement and specify the **ah**, **esp**, or **bundle** option at the `[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]` hierarchy level:

```

[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]
protocol (ah | bundle | esp);

```

Configuring the Security Parameter Index

An SPI is an arbitrary value that uniquely identifies which SA to use at the receiving host. The sending host uses the SPI to identify and select which SA to use to secure every packet. The receiving host uses the SPI to identify and select the encryption algorithm and key used to decrypt packets.



NOTE: Each manual SA must have a unique SPI and protocol combination. Use the auxiliary SPI when you configure the protocol statement to use the **bundle** option.

To configure the SPI, include the **spi** statement and specify a value (from 256 through 16,639) at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]  
spi spi-value;
```

Configuring the Auxiliary Security Parameter Index

Use the auxiliary SPI when you configure the **protocol** statement to use the **bundle** option.



NOTE: Each manual SA must have a unique SPI and protocol combination.

To configure the auxiliary SPI, include the **auxiliary-spi** statement and specify a value (from 256 through 16,639) at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]  
auxiliary-spi auxiliary-spi-value;
```

Configuring Authentication for a Manual IPsec SA

To configure an authentication algorithm, include the **authentication** statement and specify an authentication algorithm and a key at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]  
authentication {  
  algorithm (hmac-md5-96 | hmac-sha1-96 | hmac-sha-256-128)  
  key (ascii-text key | hexadecimal key);  
}
```

The algorithm can be one of the following:

- **hmac-md5-96**—Hash algorithm that authenticates packet data. It produces a 128-bit authenticator value and a 96-bit digest.
- **hmac-sha1-96**—Hash algorithm that authenticates packet data. It produces a 160-bit authenticator value and a 96-bit digest.
- **hmac-sha-256-128**—Hash algorithm that authenticates packet data. It produces a 256-bit authenticator value 256-bit digest, truncated to 128 bits.

The key can be one of the following:

- **ascii-text**—ASCII text key. With the **hmac-md5-96** option, the key contains 16 ASCII characters. With the **hmac-sha1-96** option, the key contains 20 ASCII characters.
- **hexadecimal**—Hexadecimal key. With the **hmac-md5-96** option, the key contains 32 hexadecimal characters. With the **hmac-sha1-96** option, the key contains 40 hexadecimal characters.

Configuring Encryption for a Manual IPsec SA

To configure IPsec encryption, include the **encryption** statement and specify an algorithm and key at the **[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]
  encryption {
    algorithm algorithm;
    key (ascii-text key | hexadecimal key);
  }
```

The algorithm can be one of the following:

- **des-cbc**—Encryption algorithm that has a block size of 8 bytes; its key size is 64 bits long.
- **3des-cbc**—Encryption algorithm that has a block size of 24 bytes; its key size is 192 bits long.
- **aes-128-cbc**—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- **aes-192-cbc**—Advanced Encryption Standard (AES) 192-bit encryption algorithm.
- **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption algorithm.



NOTE: For a list of Data Encryption Standard (DES) encryption algorithm weak and semiweak keys, see RFC 2409, *The Internet Key Exchange (IKE)*. The AES encryption algorithms use a software implementation that has much lower throughput, so DES remains the recommended option. For reference information on AES encryption, see RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*.

For **3des-cbc**, the first 8 bytes should differ from the second 8 bytes, and the second 8 bytes should be the same as the third 8 bytes.

If you configure an authentication proposal but do not include the encryption statement, the result is NULL encryption. Certain applications expect this result. If you configure no specific authentication or encryption values, the Junos OS uses the default values of **sha1** for the authentication and **3des-cbc** for the encryption.

The key can be one of the following:

- **ascii-text**—ASCII text key. With the **des-cbc** option, the key contains 8 ASCII characters. With the **3des-cbc** option, the key contains 24 ASCII characters.
- **hexadecimal**—Hexadecimal key. With the **des-cbc** option, the key contains 16 hexadecimal characters. With the **3des-cbc** option, the key contains 48 hexadecimal characters.



NOTE: You cannot configure encryption when you use the AH protocol.

Configuring Dynamic Security Associations

You configure dynamic SAs with a set of proposals that are negotiated by the security gateways. The keys are generated as part of the negotiation and therefore do not need to be specified in the configuration. The dynamic SA includes one or more proposals, which allow you to prioritize a list of protocols and algorithms to be negotiated with the peer.

To enable a dynamic SA, follow these steps:

1. Configure Internet Key Exchange (IKE) proposals and IKE policies associated with these proposals.
2. Configure IPsec proposals and an IPsec policy associated with these proposals.
3. Associate an SA with an IPsec policy by configuring the **dynamic** statement.

To configure a dynamic SA, include the **dynamic** statement and specify an IPsec policy name at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level. The **ike-policy** statement is optional unless you use the preshared key authentication method.

```
[edit services ipsec-vpn rule rule-name term term-name then]
dynamic {
  ike-policy policy-name;
  ipsec-policy policy-name;
}
```



NOTE: If you want to establish a dynamic SA, the attributes in at least one configured IPsec and IKE proposal must match those of its peer.

Clearing Security Associations

You can set up the router software to clear IKE or IPsec SAs automatically when the corresponding services PIC restarts or is taken offline. To configure this property, include the **clear-ike-sas-on-pic-restart** or **clear-ipsec-sas-on-pic-restart** statement at the `[edit services ipsec-vpn]` hierarchy level:

```
[edit services ipsec-vpn]
clear-ike-sas-on-pic-restart;
clear-ipsec-sas-on-pic-restart;
```

After you add this statement to the configuration, all the IKE or IPsec SAs corresponding to the tunnels in the PIC will be cleared when the PIC restarts or goes offline.

Related Documentation

- [Configuring IPsec Policies on page 420](#)
- [Configuring IPsec Proposals on page 415](#)

- [Configuring IKE Policies on page 409](#)
- [Configuring IKE Proposals on page 405](#)

Example: Configuring Manual SAs

This example shows how to create an IPsec tunnel by using manual security associations (SAs), and contains the following sections:

- [Requirements on page 391](#)
- [Overview and Topology on page 391](#)
- [Configuration on page 392](#)
- [Verification on page 403](#)

Requirements

This example uses the following hardware and software components:

- Four M Series, MX Series, or T Series routers with multiservices interfaces installed in them.
- Junos OS Release 9.4 and later.

No special configuration beyond device initialization is required before you can configure this feature.

Overview and Topology


A security association (SA) is a simplex connection that enables two hosts to securely communicate with each other by means of IPsec. There are two types of SAs: manual SA and dynamic SA. This example explains a manual SA configuration.

Manual SAs require no negotiation; all values, including the keys, are static and specified in the configuration. Manual SAs use statically defined security parameter index (SPI) values, algorithms, and keys, and require matching configurations on both ends of the tunnel. Each peer must have the same configured options for communication to take place.

Manual SAs are best suited for small, static networks where the distribution, maintenance, and tracking of keys are not difficult.

[Figure 21 on page 392](#) shows an IPsec topology that contains a group of four routers: Routers 1, 2, 3, and 4.

Figure 21: Manual SA Topology



Routers 2 and 3 establish an IPsec tunnel by using a multiservices PIC and manual SA settings. Routers 1 and 4 provide basic connectivity and are used to verify that the IPsec tunnel is operational.

Configuration

This example uses four routers, and involves the following configurations:

- Routers 1 and 4 are configured for basic OSPF connectivity with Routers 2 and 3 respectively.
- Routers 2 and 3 are configured for OSPF connectivity with Routers 1 and 4 respectively. Routers 2 and 3 are also configured to create an IPsec tunnel by using manual SAs between these two routers. To direct traffic to the IPsec tunnel through the multiservices interface, next-hop style service sets are configured on Routers 2 and 3, and the multiservices interfaces that are configured as the IPsec inside interface are added to the OSPF configuration on the respective routers.



NOTE: The interface types shown in this example are for indicative purpose only. For example, you can use `so-` interfaces instead of `ge-` and `sp-` instead of `ms-`.

This section contains:

- [Configuring Router 1 on page 392](#)
- [Configuring Router 2 on page 394](#)
- [Configuring Router 3 on page 398](#)
- [Configuring Router 4 on page 402](#)

Configuring Router 1

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI, at the [edit] hierarchy level, of Router 1.

```
set interfaces ge-1/0/1 description "to R2 ge-1/0/1"
set interfaces ge-1/0/1 unit 0 family inet address 10.1.12.2/30
```

```

set interfaces lo0 unit 0 family inet address 10.0.0.1/32
set protocols ospf area 0.0.0.0 interface ge-1/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0
set routing-options router-id 10.0.0.1

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router 1 for OSPF connectivity with Router 2:

1. Configure an Ethernet interface and loopback interface.

```

[edit interfaces]
user@router1# set ge-1/0/1 description "to R2 ge-1/0/1"
user@router1# set ge-1/0/1 unit 0 family inet address 10.1.12.2/30
user@router1# set lo0 unit 0 family inet address 10.0.0.1/32

```
2. Specify the OSPF area and associate the interfaces with the OSPF area.

```

[edit protocols]
user@router1# set ospf area 0.0.0.0 interface ge-1/0/1.0
user@router1# set ospf area 0.0.0.0 interface lo0.0

```
3. Configure the router ID.

```

[edit routing-options]
user@router1# set router-id 10.0.0.1

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration

```

user@router1# show interfaces
interfaces {
  ...
  ge-1/0/1 {
    description "to R2 ge-1/0/1";
    unit 0 {
      family inet {
        address 10.1.12.2/30;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.1/32;
      }
    }
  }
  ...
}

user@router1# show protocols ospf
ospf {

```

```

    area 0.0.0.0 {
        interface ge-1/0/1.0;
        interface lo0.0;
    }
}

user@router1# show routing-options
routing-options {
    router-id 10.0.0.1;
}

```

Configuring Router 2

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI, at the [edit] hierarchy level, of Router 2.

Configuring Interfaces and OSPF Connectivity (with Router 1 and Router 3) on Router 2

```

set interfaces ge-1/0/0 unit 0 description "to R3 ge-1/0/0"
set interfaces ge-1/0/0 unit 0 family inet address 10.1.15.1/30
set interfaces ge-1/0/1 unit 0 description "to R1 ge-1/0/0"
set interfaces ge-1/0/1 unit 0 family inet address 10.1.12.1/30
set interfaces ms-1/2/0 unit 0 family inet
set interfaces ms-1/2/0 unit 1 family inet
set interfaces ms-1/2/0 unit 1 service-domain inside
set interfaces ms-1/2/0 unit 2 family inet
set interfaces ms-1/2/0 unit 2 service-domain outside
set interfaces lo0 unit 0 family inet address 10.0.0.2/32
set protocols ospf area 0.0.0.0 interface ge-1/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ms-1/2/0.1
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  remote-gateway 10.1.15.2
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional protocol esp
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional spi 261
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional authentication algorithm hmac-sha1-96
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional authentication key ascii-text demokeyipsecmanualsa
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional encryption algorithm des-cbc
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional encryption key ascii-text manualsa
set services ipsec-vpn rule demo-rule-r1-manual-sa match-direction input
set services service-set demo-ss-manual-sa next-hop-service inside-service-interface
  ms-1/2/0.1
set services service-set demo-ss-manual-sa next-hop-service outside-service-interface
  ms-1/2/0.2
set services service-set demo-ss-manual-sa ipsec-vpn-options local-gateway 10.1.15.1
set services service-set demo-ss-manual-sa ipsec-vpn-rules demo-rule-r1-manual-sa

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure OSPF connectivity and IPsec tunnel parameters on Router 2:

1. Configure interface properties. In this step, you configure two Ethernet interfaces (ge-1/0/0 and ge-1/0/1), a loopback interface, and a multiservices interface (ms-1/2/0).

```
[edit interfaces]
user@router2# set ge-1/0/0 unit 0 description "to R3 ge-1/0/0"
user@router2# set ge-1/0/0 unit 0 family inet address 10.1.15.1/30
user@router2# set ge-1/0/1 unit 0 description "to R1 ge-1/0/0"
user@router2# set ge-1/0/1 unit 0 family inet address 10.1.12.1/30
user@router2# set ms-1/2/0 unit 0 family inet
user@router2# set ms-1/2/0 unit 1 family inet
user@router2# set ms-1/2/0 unit 1 service-domain inside
user@router2# set ms-1/2/0 unit 2 family inet
user@router2# set ms-1/2/0 unit 2 service-domain outside
user@router2# set lo0 unit 0 family inet address 10.0.0.2/32
```

2. Specify the OSPF area and associate the interfaces with the OSPF area.

```
[edit protocols]
user@router2# set ospf area 0.0.0.0 interface ge-1/0/1.0
user@router2# set ospf area 0.0.0.0 interface lo0.0
user@router2# set ospf area 0.0.0.0 interface ms-1/2/0.1
```

3. Configure the router ID.

```
[edit routing-options]
user@router2# set router-ID 10.0.0.2
```

4. Configure an IPsec rule. In this step, you configure an IPsec rule and specify manual SA parameters, such as the remote-gateway address, authentication and encryption properties, and so on.

```
[edit services ipsec-vpn]
user@router2# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  remote-gateway 10.1.15.2
user@router2# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional protocol esp
user@router2# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional spi 261
user@router2# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional authentication algorithm hmac-sha1-96
user@router2# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional authentication key ascii-text
  demokeyipsecmanualsa
user@router2# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional encryption algorithm des-cbc
user@router2# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional encryption key ascii-text manualsa
user@router2# set rule demo-rule-r1-manual-sa match-direction input
```

5. Configure a next-hop style service set, specify the local-gateway address, and associate the IPsec VPN rule with the service set.

```
[edit services]
user@router2# set service-set demo-ss-manual-sa next-hop-service
    inside-service-interface ms-1/2/0.1
user@router2# set service-set demo-ss-manual-sa next-hop-service
    outside-service-interface ms-1/2/0.2
user@router2# set service-set demo-ss-manual-sa ipsec-vpn-options local-gateway
    10.1.15.1
user@router2# set service-set demo-ss-manual-sa ipsec-vpn-rules
    demo-rule-r1-manual-sa
```

6. Commit the configuration.

```
[edit]
user@router2# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show routing-options**, and **show services** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration

```
user@router1# show interfaces
interfaces {
  ...
  ge-1/0/0 {
    unit 0 {
      description "to R3 ge-1/0/0";
      family inet {
        address 10.1.15.1/30;
      }
    }
  }
  ge-1/0/1 {
    unit 0 {
      description "to R1 ge-1/0/1";
      family inet {
        address 10.1.12.1/30;
      }
    }
  }
  ms-1/2/0 {
    unit 0 {
      family inet;
    }
    unit 1 {
      family inet;
      service-domain inside;
    }
    unit 2 {
      family inet;
      service-domain outside;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.2/32;
      }
    }
  }
}
```

```

    }
  }
}
...
}

user@router2# show protocols ospf
protocols {
  ospf {
    area 0.0.0.0 {
      interfaces ge-1/0/1.0;
      interface lo0;
      interface ms-1/2/0;
    }
  }
}

user@router2# show routing-options
routing-options {
  router-id 10.0.0.2;
}

user@router2# show services
services {
  ipsec-vpn {
    rule demo-rule-r1-manual-sa {
      term demo-term-manual-sa {
        then {
          remote-gateway 10.1.15.2;
          manual {
            direction bidirectional {
              protocol esp;
              spi 261;
              authentication {
                algorithm hmac-sha1-96;
                key ascii-text
                  "$9$km5FctOcyKn/yKM8dVqmf5QntpBcyKturWLVbz369pBIRSM87revLX-2g";
                ## SECRET-DATA
              }
              encryption {
                algorithm des-cbc;
                key ascii-text "$9$n2Hi/tO1lcvWxylK8LNY2Tz36/t"; ## SECRET-DATA
              }
            }
          }
        }
      }
    }
    match-direction input;
  }
}

service-set demo-ss-manual-sa {
  next-hop-service {
    inside-service-interface ms-1/2/0.1;
    outside-service-interface ms-1/2/0.2;
  }
  ipsec-vpn-options {
    local-gateway 10.1.15.1;
  }
}

```

```
    }  
    ipsec-vpn-rules demo-rule-r1-manual-sa;  
  }  
}
```

Configuring Router 3

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI, at the [edit] hierarchy level, of Router 3.

```
set interfaces ge-1/0/1 unit 0 description "to R4 ge-1/0/1"  
set interfaces ge-1/0/0 unit 0 family inet address 10.1.56.1/30  
set interfaces ge-1/0/0 unit 0 description "to R2 ge-1/0/0"  
set interfaces ge-1/0/0 unit 0 family inet address 10.1.15.2/30  
set interfaces ms-1/2/0 unit 0 family inet  
set interfaces ms-1/2/0 unit 1 family inet  
set interfaces ms-1/2/0 unit 1 service-domain inside  
set interfaces ms-1/2/0 unit 2 family inet  
set interfaces ms-1/2/0 unit 2 service-domain outside  
set interfaces lo0 unit 0 family inet address 10.0.0.3/32  
set protocols ospf area 0.0.0.0 interface ge-1/0/1.0  
set protocols ospf area 0.0.0.0 interface lo0.0  
set protocols ospf area 0.0.0.0 interface ms-1/2/0.1  
set routing-options router-id 10.0.0.3  
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then  
  remote-gateway 10.1.15.1  
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then  
  manual direction bidirectional protocol esp  
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then  
  manual direction bidirectional spi 261  
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then  
  manual direction bidirectional authentication algorithm hmac-sha1-96  
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then  
  manual direction bidirectional authentication key ascii-text demokeyipsecmanualsa  
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then  
  manual direction bidirectional encryption algorithm des-cbc  
set services ipsec-vpn rule demo-rule-r1-manual-sa term demo-term-manual-sa then  
  manual direction bidirectional encryption key ascii-text manualsa  
set services ipsec-vpn rule demo-rule-r1-manual-sa match-direction input  
set services service-set demo-ss-manual-sa next-hop-service inside-service-interface  
  ms-1/2/0.1  
set services service-set demo-ss-manual-sa next-hop-service outside-service-interface  
  ms-1/2/0.2  
set services service-set demo-ss-manual-sa ipsec-vpn-options local-gateway 10.1.15.2  
set services service-set demo-ss-manual-sa ipsec-vpn-rules demo-rule-r1-manual-sa
```


Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure OSPF connectivity and IPsec tunnel parameters on Router 3:

1. Configure interface properties. In this step, you configure two Ethernet interfaces (ge-1/0/0 and ge-1/0/1), a loopback interface, and a multiservices interface (ms-1/2/0).

```
[edit interfaces]
user@router3# set ge-1/0/0 unit 0 description "to R4 ge-1/0/0"
user@router3# set ge-1/0/0 unit 0 family inet address 10.1.56.1/30
user@router3# set ge-1/0/1 unit 0 description "to R2 ge-1/0/1"
user@router3# set ge-1/0/1 unit 0 family inet address 10.1.15.2/30
user@router3# set ms-1/2/0 unit 0 family inet
user@router3# set ms-1/2/0 unit 1 family inet
user@router3# set ms-1/2/0 unit 1 service-domain inside
user@router3# set ms-1/2/0 unit 2 family inet
user@router3# set ms-1/2/0 unit 2 service-domain outside
user@router3# set lo0 unit 0 family inet address 10.0.0.3/32
```

2. Specify the OSPF area and associate the interfaces with the OSPF area.

```
[edit protocols]
user@router3# set ospf area 0.0.0.0 interface ge-1/0/1.0
user@router3# set ospf area 0.0.0.0 interface lo0.0
user@router3# set ospf area 0.0.0.0 interface ms-1/2/0.1
```

3. Configure a router ID.

```
[edit routing-options]
user@router3# set router-id 10.0.0.3
```

4. Configure an IPsec rule. In this step, you configure an IPsec rule and specify manual SA parameters, such as the remote-gateway address, authentication and encryption properties, and so on.

```
[edit services ipsec-vpn]
user@router3# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  remote-gateway 10.1.15.1
user@router3# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional protocol esp
user@router3# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional spi 261
user@router3# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional authentication algorithm hmac-sha1-96
user@router3# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional authentication key ascii-text
  demokeyipsecmanualsa
user@router3# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional encryption algorithm des-cbc
user@router3# set rule demo-rule-r1-manual-sa term demo-term-manual-sa then
  manual direction bidirectional encryption key ascii-text manualsa
user@router3# set rule demo-rule-r1-manual-sa match-direction input
```

5. Configure a next-hop style service set, specify the local-gateway address, and associate the IPsec VPN rule with the service set.

```
[edit services]
user@router3# set service-set demo-ss-manual-sa next-hop-service
    inside-service-interface ms-1/2/0.1
user@router3# set service-set demo-ss-manual-sa next-hop-service
    outside-service-interface ms-1/2/0.2
user@router3# set service-set demo-ss-manual-sa ipsec-vpn-options local-gateway
    10.1.15.2
user@router3# set service-set demo-ss-manual-sa ipsec-vpn-rules
    demo-rule-r1-manual-sa
```

6. Commit the configuration.

```
[edit]
user@router3# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show routing-options**, and **show services** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration

```
user@router3# show interfaces
interfaces {
  ge-1/0/1 {
    unit 0 {
      description "to R4 ge-1/0/1";
      family inet {
        address 10.1.56.1/30;
      }
    }
  }
  ge-1/0/0 {
    unit 0 {
      description "to R2 ge-1/0/0";
      family inet {
        address 10.1.15.2/30;
      }
    }
  }
  ms-1/2/0 {
    unit 0 {
      family inet;
    }
    unit 1 {
      family inet;
      service-domain inside;
    }
    unit 2 {
      family inet;
      service-domain outside;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.3/32;
      }
    }
  }
}
```

```

    }
  }
}

user@router3# show protocols ospf
protocols {
  ospf {
    area 0.0.0.0 {
      interface ge-1/0/1.0;
      interface lo0.0;
      interface ms-1/2/0.1;
    }
  }
}

user@router3# show routing-options
routing-options {
  router-id 10.0.0.3;
}

user@router3# show services
services {
  ipsec-vpn {
    rule demo-rule-r1-manual-sa {
      term demo-term-manual-sa {
        then {
          remote-gateway 10.1.15.1;
          manual {
            direction bidirectional {
              protocol esp;
              spi 261;
              authentication {
                algorithm hmac-sha1-96;
                key ascii-text
                  "$9$km5FCtOcyKn/yKM8dVqmf5QntpBcyKturvWLVbz369pBIRSM87revLX-2g";
                ## SECRET-DATA
              }
              encryption {
                algorithm des-cbc;
                key ascii-text "$9$n2Hi/tO1lcvWxykK8LNY2Tz36/t"; ## SECRET-DATA
              }
            }
          }
        }
      }
    }
    match-direction input;
  }
}

service-set demo-ss-manual-sa {
  next-hop-service {
    inside-service-interface ms-1/2/0.1;
    outside-service-interface ms-1/2/0.2;
  }
  ipsec-vpn-options {
    local-gateway 10.1.15.2;
  }
  ipsec-vpn-rules demo-rule-r1-manual-sa;
}

```

```
}  
}
```

Configuring Router 4

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI, at the [edit] hierarchy level, of Router 4.

```
set interfaces ge-1/0/1 description "to R3 ge-1/0/1"  
set interfaces ge-1/0/1 unit 0 family inet address 10.1.56.2/30  
set interfaces lo0 unit 0 family inet address 10.0.0.4/32  
set protocols ospf area 0.0.0.0 interface ge-1/0/1.0  
set protocols ospf area 0.0.0.0 interface lo0.0  
set routing-options router-id 10.0.0.4
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To set up OSPF connectivity with Router 3

1. Configure the interfaces. In this step, you configure an Ethernet interface (ge-1/0/1) and a loopback interface.

```
user@router4# set interfaces ge-1/0/1 description "to R3 ge-1/0/1"  
user@router4# set interfaces ge-1/0/1 unit 0 family inet address 10.1.56.2/30  
user@router4# set interfaces lo0 unit 0 family inet address 10.0.0.4/32
```

2. Specify the OSPF area and associate the interfaces with the OSPF area.

```
user@router4# set protocols ospf area 0.0.0.0 interface ge-1/0/1.0  
user@router4# set protocols ospf area 0.0.0.0 interface lo0.0
```

3. Configure the router ID.

```
[edit routing-options]  
user@router4# set router-id 10.0.0.4
```

4. Commit the configuration.

```
[edit]  
user@router4# commit
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration

```
user@router4# show interfaces  
interfaces {  
  ge-1/0/1 {  
    description "to R3 ge-1/0/1";  
    unit 0 {  
      family inet {
```

```

        address 10.1.56.2/30;
    }
}
lo0{
    unit 0 {
        family inet {
            address 10.0.0.4/32;
        }
    }
}
}

user@router4# show routing-options
routing-options {
    router-id 10.0.0.4;
}

user@router4# show protocols ospf
protocols {
    ospf {
        area 0.0.0.0 {
            interface lo0.0;
            interface ge-1/0/1.0;
        }
    }
}

```

Verification

To confirm that the manual SA configuration is working properly, perform the following tasks:

- [Verifying Traffic Flow Through the IPsec Tunnel on page 403](#)
- [Verifying the Security Associations on Router 2 on page 404](#)
- [Verifying the Security Associations on Router 3 on page 404](#)

Verifying Traffic Flow Through the IPsec Tunnel

Purpose Verify that the IPsec tunnel carries traffic between Router 1 and Router 4.

Action Issue a **ping** command from Router 1 to **lo0** on Router 4.

```

user@router1> ping 10.0.0.4
PING 10.0.0.4 (10.0.0.4): 56 data bytes
64 bytes from 10.0.0.4: icmp_seq=0 ttl=254 time=1.375 ms
64 bytes from 10.0.0.4: icmp_seq=1 ttl=254 time=18.375 ms
64 bytes from 10.0.0.4: icmp_seq=2 ttl=254 time=1.120 ms
^C
--- 10.0.0.4 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.120/6.957/18.375/8.075 ms

```

Meaning The output shows that Router 1 is able to reach Router 4 over the IPsec tunnel.

Verifying the Security Associations on Router 2

Purpose Verify that the security associations are active on Router 2 and that the traffic is flowing over the IPsec tunnel.

Action • To verify that the security associations are active, Issue **show services ipsec-vpn ipsec security-associations detail** on Router 2.

```
user@router2> show services ipsec-vpn ipsec security-associations detail
Service set: demo-ss-manual-sa
Rule: demo-rule-r1-manual-sa, Term: demo-term-manual-sa,
Tunnel index: 1
Local gateway: 10.1.15.1, Remote gateway: 10.1.15.2
Local identity: ipv4_subnet(any:0,[0..7]=10.0.0.0/8)
Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Direction: inbound, SPI: 261, AUX-SPI: 0
Mode: tunnel, Type: manual, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
Anti-replay service: Disabled
Direction: outbound, SPI: 261, AUX-SPI: 0
Mode: tunnel, Type: manual, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
Anti-replay service: Disabled
```

• To verify that traffic is traveling over the bidirectional IPsec tunnel, issue **show services ipsec-vpn ipsec statistics** on Router 2.

```
user@router2# show services ipsec-vpn ipsec statistics
PIC: ms-1/2/0, Service set: demo-ss-manual-sa
sESP Statistics:
Encrypted bytes: 1616
Decrypted bytes: 1560
Encrypted packets: 20
Decrypted packets: 19
AH Statistics:
Input bytes: 0
Output bytes: 0
Input packets: 0
Output packets: 0
Errors:
AH authentication failures: 0, Replay errors: 0
ESP authentication failures: 0, ESP decryption failures: 0
Bad headers: 0, Bad trailers: 0
```

Meaning The **show services ipsec-vpn ipsec security-associations detail** command output shows the SA properties that you configured.

The **show services ipsec-vpn ipsec statistics** command output shows the traffic flow over the IPsec tunnel.

Verifying the Security Associations on Router 3

Purpose Verify the security associations and flow of traffic over the IPsec tunnel.

Action • To verify that the security associations are active, Issue **show services ipsec-vpn ipsec security-associations detail** on Router 3.

```

user@router3> show services ipsec-vpn ipsec security-associations detail
Service set: demo-ss-manual-sa
Rule: demo-rule-r1-manual-sa, Term: demo-term-manual-sa,
Tunnel index: 1
Local gateway: 10.1.15.2, Remote gateway: 10.1.15.1
Local identity: ipv4_subnet(any:0,[0..7]=10.0.0.0/8)
Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Direction: inbound, SPI: 261, AUX-SPI: 0
Mode: tunnel, Type: manual, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
Anti-replay service: Disabled
Direction: outbound, SPI: 261, AUX-SPI: 0
Mode: tunnel, Type: manual, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
Anti-replay service: Disabled

```

- To verify that traffic is traveling over the bidirectional IPsec tunnel, issue **show services ipsec-vpn ipsec statistics** on Router 3.

```

user@router3# show services ipsec-vpn ipsec statistics
PIC: ms-1/2/0, Service set: demo-ss-manual-sa
ESP Statistics:
Encrypted bytes: 1560
Decrypted bytes: 1616
Encrypted packets: 19
Decrypted packets: 20
AH Statistics:
Input bytes: 0
Output bytes: 0
Input packets: 0
Output packets: 0
Errors:
AH authentication failures: 0, Replay errors: 0
ESP authentication failures: 0, ESP decryption failures: 0
Bad headers: 0, Bad trailers: 0

```

Meaning The **show services ipsec-vpn ipsec security-associations detail** command output shows the SA properties that you configured.

The **show services ipsec-vpn ipsec statistics** command output shows the traffic flow over the IPsec tunnel.

- Related Documentation**
- [Understanding Junos VPN Site Secure on page 371](#)
 - [Configuring Security Associations on page 385](#)
 - [Example: Configuring IKE Dynamic SAs on page 466](#)

Configuring IKE Proposals

Dynamic security associations (SAs) require IKE configuration. With dynamic SAs, you configure IKE first, and then the SA. IKE creates the dynamic SAs and negotiates them for IPsec. The IKE configuration defines the algorithms and keys used to establish the secure IKE connection with the peer security gateway.

You can configure one or more IKE proposals. Each proposal is a list of IKE attributes to protect the IKE connection between the IKE host and its peer.

To configure an IKE proposal, include the **proposal** statement and specify a name at the **[edit services ipsec-vpn ike]** hierarchy level:

```
[edit services ipsec-vpn ike]
proposal proposal-name {
  authentication-algorithm (md5 | sha1 | sha-256);
  authentication-method (pre-shared-key | rsa-signatures);
  dh-group (group1 | group2 | group5 | group14 | group19 | group20);
  encryption-algorithm algorithm;
  lifetime-seconds seconds;
}
```

This section includes the following topics:

- [Configuring the Authentication Algorithm for an IKE Proposal on page 406](#)
- [Configuring the Authentication Method for an IKE Proposal on page 406](#)
- [Configuring the Diffie-Hellman Group for an IKE Proposal on page 407](#)
- [Configuring the Encryption Algorithm for an IKE Proposal on page 408](#)
- [Configuring the Lifetime for an IKE SA on page 408](#)
- [Example: Configuring an IKE Proposal on page 409](#)

Configuring the Authentication Algorithm for an IKE Proposal

To configure the authentication algorithm for an IKE proposal, include the **authentication-algorithm** statement at the **[edit services ipsec-vpn ike proposal proposal-name]** hierarchy level:

```
[edit services ipsec-vpn ike proposal proposal-name]
authentication-algorithm (md5 | sha1 | sha-256);
```

The authentication algorithm can be one of the following:

- **md5**—Produces a 128-bit digest.
- **sha1**—Produces a 160-bit digest.
- **sha-256**—Produces a 256-bit digest.



NOTE: For reference information on Secure Hash Algorithms (SHAs), see Internet draft [draft-eastlake-sha2-02.txt](#), *Secure Hash Algorithms (SHA and HMAC-SHA)* (expires July 2006).

Configuring the Authentication Method for an IKE Proposal

To configure the authentication method for an IKE proposal, include the **authentication-method** statement at the **[edit services ipsec-vpn ike proposal proposal-name]** hierarchy level:


```
[edit services ipsec-vpn ike proposal proposal-name]
authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
```



NOTE: In IKEv1, the authentication method for SAs is negotiated with the remote peer based on the type of authentication method configured in the IKE proposal. In IKEv2, such a negotiation is not performed with the remote peer. Instead, each IKE peer uses the authentication method that is locally configured for them.

For SAs in IKEv2, the authentication method is the default value as IKEv1 if an authentication method is not configured in the IKE proposal. If you are configuring an authentication method for IKEv2, you must have the same authentication method configured for all proposals referenced in the policy.

The authentication method can be one of the following:

- **pre-shared-keys**—A key derived from an out-of-band mechanism; the key authenticates the exchanges
- **rsa-signatures**—Public key algorithm (supports encryption and digital signatures)

Configuring the Diffie-Hellman Group for an IKE Proposal

Diffie-Hellman is a public-key cryptography scheme that allows two parties to establish a shared secret over an insecure communications channel. It is also used within IKE to establish session keys.

To configure the Diffie-Hellman group for an IKE proposal, include the **dh-group** statement at the `[edit services ipsec-vpn ike proposal proposal-name]` hierarchy level:

```
[edit services ipsec-vpn ike proposal proposal-name]
dh-group (group1 | group2 | group5 | group14 | group19 | group20);
```

The group can be one of the following:

- **group1**—Specifies that IKE uses the 768-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- **group2**—Specifies that IKE uses the 1024-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- **group5**—Specifies that IKE uses the 1536-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- **group14**—Specifies that IKE uses the 2048-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- **group19**—Specifies that IKE uses the 256-bit random Elliptic Curve Diffie-Hellman Group when performing the new Diffie-Hellman exchange.
- **group20**—Specifies that IKE uses the 384-bit random Elliptic Curve Diffie-Hellman Group when performing the new Diffie-Hellman exchange.

Using a Diffie-Hellman group based on a greater number of bits results a more secure IKE tunnel than using a group based on fewer bits. However, this additional security might require additional processing time.

Configuring the Encryption Algorithm for an IKE Proposal

To configure the encryption algorithm for an IKE proposal, include the **encryption-algorithm** statement at the **[edit services ipsec-vpn ike proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ike proposal proposal-name]  
  encryption-algorithm algorithm;
```

The encryption algorithm can be one of the following:

- **3des-cbc**—Cipher block chaining encryption algorithm with a key size of 24 bytes; its key size is 192 bits long.
- **des-cbc**—Cipher block chaining encryption algorithm with a key size of 8 bytes; its key size is 56 bits long.
- **aes-128-cbc**—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- **aes-192-cbc**—Advanced Encryption Standard (AES) 192-bit encryption algorithm.
- **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption algorithm.



NOTE: For a list of Data Encryption Standard (DES) encryption algorithm weak and semiweak keys, see RFC 2409, *The Internet Key Exchange (IKE)*. The AES encryption algorithms use a software implementation that has much lower throughput, so DES remains the recommended option.

For 3des-cbc, the first 8 bytes should differ from the second 8 bytes, and the second 8 bytes should be the same as the third 8 bytes.

If you configure an authentication proposal but do not include the encryption statement, the result is NULL encryption. Certain applications expect this result. If you configure no specific authentication or encryption values, the Junos OS uses the default values of sha1 for the authentication and 3des-cbc for the encryption.

Configuring the Lifetime for an IKE SA

The **lifetime-seconds** statement sets the lifetime of an IKE SA. When the IKE SA expires, it is replaced by a new SA (and SPI) or the IPsec connection is terminated.

To configure the lifetime for an IKE SA, include the **lifetime-seconds** statement at the **[edit services ipsec-vpn ike proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ike proposal proposal-name]  
  lifetime-seconds seconds;
```

By default, the IKE SA lifetime is 3600 seconds. The range is from 180 through 86,400 seconds.



NOTE: In IKEv1, the lifetime for SAs is negotiated with the remote peer based on the type of lifetime configured in the IKE proposal. In IKEv2, such a negotiation is not performed with the remote peer. Instead, each IKE peer uses the lifetime that is locally configured for them.

For SAs in IKEv2, the lifetime is either the default value as IKEv1 (if another lifetime is not configured in the IKE proposal) or all IKEv2 proposals in the IKE policy must be configured with the same lifetime value.



NOTE: For IKE proposals, there is only one SA lifetime value, specified by the Junos OS. IPsec proposals use a different mechanism.

Example: Configuring an IKE Proposal

Configure an IKE proposal:

```
[edit services ipsec-vpn ike]
proposal ike-proposal {
  authentication-method pre-shared-keys;
  dh-group group1;
  authentication-algorithm sha1;
  encryption-algorithm 3des-cbc;
}
```

Related Documentation

- [Configuring IPsec Proposals on page 415](#)
- [Configuring IKE Policies on page 409](#)
- [Configuring IPsec Policies on page 420](#)
- [Configuring Security Associations on page 385](#)

Configuring IKE Policies

An IKE policy defines a combination of security parameters (IKE proposals) to be used during IKE negotiation. It defines a peer address and the proposals needed for that connection. Depending on which authentication method is used, it defines the preshared key for the given peer or the local certificate. During the IKE negotiation, IKE looks for an IKE policy that is the same on both peers. The peer that initiates the negotiation sends all its policies to the remote peer, and the remote peer tries to find a match.

A match is made when both policies from the two peers have a proposal that contains the same configured attributes. If the lifetimes are not identical, the shorter lifetime between the two policies (from the host and peer) is used. The configured preshared key must also match its peer.

Starting with Junos OS Release 11.4, both IKEv1 and IKEv2 are supported by default on all M Series, MX Series, and T Series routers. You can configure the specific IKE phase to be supported for the negotiation. However, if only IKEv1 is supported, the Junos OS rejects

IKEv2 negotiations. Similarly, if only IKEv2 is supported, the Junos OS rejects all IKEv1 negotiations.

The key management process (kmd) daemon determines which version of IKE is used in a negotiation. If kmd is the IKE initiator, it uses IKEv1 by default and retains the configured version for negotiations. If kmd is the IKE responder, it accepts connections from both IKEv1 and IKEv2.

You can create multiple, prioritized proposals at each peer to ensure that at least one proposal matches a remote peer's proposal.

First, you configure one or more IKE proposals; then you associate these proposals with an IKE policy. You can also prioritize a list of proposals used by IKE in the **policy** statement by listing the proposals you want to use, from first to last.

To configure an IKE policy, include the **policy** statement and specify a policy name at the **[edit services ipsec-vpn ike]** hierarchy level:

```
[edit services ipsec-vpn ike]
policy policy-name {
  description description;
  local-certificate identifier;
  local-id (ipv4_addr ipv4-address | ipv6_addr ipv6-address | key-id identifier);
  version (1 | 2);
  mode (aggressive | main);
  pre-shared-key (ascii-text key | hexadecimal key);
  proposals [ proposal-names ];
  remote-id {
    any-remote-id;
    ipv4_addr [ values ];
    ipv6_addr [ values ];
    key_id [ values ];
  }
  respond-bad-spi max-responses;
}
```

This section includes the following topics:

- [Configuring the IKE Phase on page 411](#)
- [Configuring the Mode for an IKE Policy on page 411](#)
- [Configuring the Proposals in an IKE Policy on page 411](#)
- [Configuring the Preshared Key for an IKE Policy on page 411](#)
- [Configuring the Local Certificate for an IKE Policy on page 412](#)
- [Configuring the Description for an IKE Policy on page 413](#)
- [Configuring Local and Remote IDs for IKE Phase 1 Negotiation on page 413](#)
- [Enabling Invalid SPI Recovery on page 414](#)
- [Example: Configuring an IKE Policy on page 414](#)

Configuring the IKE Phase

Starting with Junos OS Release 11.4, both IKEv1 and IKEv2 are supported by default on all M Series, MX Series, and T Series routers. You can configure the specific IKE phase to be supported for the negotiation. However, if only IKEv1 is supported, the Junos OS rejects IKEv2 negotiations. Similarly, if only IKEv2 is supported, the Junos OS rejects all IKEv1 negotiations.

To configure the IKE phase used, include the **version** statement at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]  
  version (1 | 2);
```

Configuring the Mode for an IKE Policy

IKE policy has two modes: aggressive and main. By default, main mode is enabled. Main mode uses six messages, in three exchanges, to establish the IKE SA. (These three steps are IKE SA negotiation, a Diffie-Hellman exchange, and authentication of the peer.) Main mode also allows a peer to hide its identity.

Aggressive mode also establishes an authenticated IKE SA and keys. However, aggressive mode uses half the number of messages, has less negotiation power, and does not provide identity protection. The peer can use the aggressive or main mode to start IKE negotiation; the remote peer accepts the mode sent by the peer.



NOTE: The mode configuration is required only if the **version** option is set to 1.

To configure the mode for an IKE policy, include the **mode** statement and specify **aggressive** or **main** at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]  
  mode (aggressive | main);
```

Configuring the Proposals in an IKE Policy

The IKE policy includes a list of one or more proposals associated with an IKE policy.

To configure the proposals in an IKE policy, include the **proposals** statement and specify one or more proposal names at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]  
  proposals [ proposal-names ];
```

Configuring the Preshared Key for an IKE Policy

When you include the **authentication-method pre-shared-keys** statement at the **[edit services ipsec-vpn ike proposal *proposal-name*]** hierarchy level, IKE policy preshared keys authenticate peers. You must manually configure a preshared key, which must match

that of its peer. The preshared key can be an ASCII text (alphanumeric) key or a hexadecimal key.

To configure the preshared key in an IKE policy, include the **pre-shared-key** statement and a key at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]  
pre-shared-key (ascii-text key | hexadecimal key);
```

The key can be one of the following:

- **ascii-text**—ASCII text key. With the **des-cbc** option, the key contains 8 ASCII characters. With the **3des-cbc** option, the key contains 24 ASCII characters.
- **hexadecimal**—Hexadecimal key. With the **des-cbc** option, the key contains 16 hexadecimal characters. With the **3des-cbc** option, the key contains 48 hexadecimal characters.

Configuring the Local Certificate for an IKE Policy

When you include the **authentication-method rsa-signatures** statement at the **[edit services ipsec-vpn ike proposal *proposal-name*]** hierarchy level, public key infrastructure (PKI) digital certificates authenticate peers. You must identify a local certificate that is sent to the peer during the IKE authentication phase.

To configure the local certificate for an IKE policy, include the **local-certificate** statement at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]  
local-certificate identifier;
```

The **local-certificate** statement specifies the identifier used to obtain the end entity's certificate from the certification authority. Configuring it in an IKE policy allows you the flexibility of using a separate certificate with each remote peer if that is needed. You must also specify the identity of the certification authority by configuring the **ca-profile** statement at the **[edit security pki]** hierarchy level.

You can use the configured profiles to establish a set of trusted certification authorities for use with a particular service set. This enables you to configure separate service sets for individual clients to whom you are providing IP services; the distinct service sets provide logical separation of one set of IKE sessions from another, using different local gateway addresses, or *virtualization*. To configure the set of trusted certification authorities, include the **trusted-ca** statement at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level:

```
[edit services service-set service-set-name ipsec-vpn-options]  
trusted-ca ca-profile;
```

See the following to configure a certificate revocation list:

- [Configuring a Certificate Revocation List on page 413](#)

Configuring a Certificate Revocation List

A certificate revocation list (CRL) contains a list of digital certificates that have been cancelled before their expiration date. When a participating peer uses a digital certificate, it checks the certificate signature and validity. It also acquires the most recently issued CRL and checks that the certificate serial number is not on that CRL.



NOTE: By default, certificate revocation list verification is enabled. You can disable CRL verification by including the `disable` statement at the `[edit security pki ca-profile ca-profile-name revocation-check]` hierarchy level.

By default, if the router either cannot access the Lightweight Directory Access Protocol (LDAP) URL or retrieve a valid certificate revocation list, certificate verification fails and the IPsec tunnel is not established. To override this behavior and permit the authentication of the IPsec peer when the CRL is not downloaded, include the `disable on-download-failure` statement at the `[edit security pki ca-profile ca-profile-name revocation-check crl]` hierarchy level.

To use the CA certificate revocation list, you include statements at the `[edit security pki ca-profile ca-profile-name revocation-check]` hierarchy level. For details, see the [Junos OS System Basics Configuration Guide](#).

Configuring the Description for an IKE Policy

To specify an optional text description for an IKE policy, include the `description` statement at the `[edit services ipsec-vpn ike policy policy-name]` hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]
description description;
```

Configuring Local and Remote IDs for IKE Phase 1 Negotiation

You can optionally specify local identifiers for use in IKE phase 1 negotiation. If the `local-id` statement is omitted, the local gateway address is used.

To specify one or more local IDs, include the `local-id` statement at the `[edit services ipsec-vpn ike policy policy-name]` hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]
local-id (ipv4_addr ipv4-address | ipv6_addr ipv6-address | key-id identifier);
```

You can also specify remote gateway identifiers for which the IKE policy is used. The remote gateway address in which this policy is defined is added by default.

To specify one or more remote IDs, include the `remote-id` statement at the `[edit services ipsec-vpn ike policy policy-name]` hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]
remote-id {
  any-remote-id;
  ipv4_addr [ values ];
  ipv6_addr [ values ];
```

```
key_id [ values ];  
}
```

The **any-remote-id** option allows any remote address to connect. This option is supported only in dynamic endpoints configurations and cannot be configured along with specific values.

Enabling Invalid SPI Recovery

When peers in a security association (SA) become unsynchronized, packets with invalid security parameter index (SPI) values can be sent out, and the receiving peer drops these packets. For example, this could occur when one of the peers reboots. You can enable the device to recover when packets with invalid SPIs are received by resynchronizing the SAs.

To enable recovery from invalid SPI values, include the **respond-bad-spi** statement at the **[edit services ipsec-vpn ike policy] *policy-name*** hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]  
respond-bad-spi max-responses;
```

Example: Configuring an IKE Policy

Define two IKE policies: **policy 10.1.1.2** and **policy 10.1.1.1**. Each policy is associated with **proposal-1** and **proposal-2**. The following configuration uses only IKEv1 for negotiation.

```
[edit services ipsec-vpn]  
ike {  
  proposal proposal-1 {  
    authentication-method pre-shared-keys;  
    dh-group group1;  
    authentication-algorithm sha1;  
    encryption-algorithm 3des-cbc;  
    lifetime-seconds 1000;  
  }  
  proposal proposal-2 {  
    authentication-method pre-shared-keys;  
    dh-group group2;  
    authentication-algorithm md5;  
    encryption-algorithm des-cbc;  
    lifetime-seconds 10000;  
  }  
  proposal proposal-3 {  
    authentication-method rsa-signatures;  
    dh-group group2;  
    authentication-algorithm md5;  
    encryption-algorithm des-cbc;  
    lifetime-seconds 10000;  
  }  
  policy 10.1.1.2 {  
    mode main;  
    proposals [ proposal-1 proposal-2 ];  
    pre-shared-key ascii-text example-pre-shared-key;  
  }  
  policy 10.1.1.1 {
```



```

local-certificate certificate-file-name;
local-key-pair private-public-key-file;
mode aggressive;
proposals [ proposal-2 proposal-3 ]
pre-shared-key hexadecimal 0102030abbcdd;
}

```



NOTE: Updates to the current IKE proposal and policy configuration are not applied to the current IKE SA; updates are applied to new IKE SAs.

If you want the new updates to take immediate effect, you must clear the existing IKE security associations so that they will be reestablished with the changed configuration. For information about how to clear the current IKE security association, see [clear services ipsec-vpn ike security-associations](#).

Related Documentation

- [Configuring Dynamic Endpoints for IPsec Tunnels on page 455](#)
- [Configuring IKE Proposals on page 405](#)
- [Configuring IPsec Policies on page 420](#)
- [Configuring IPsec Proposals on page 415](#)
- [Configuring Security Associations on page 385](#)

Configuring IPsec Proposals

An IPsec proposal lists protocols and algorithms (security services) to be negotiated with the remote IPsec peer.

To configure an IPsec proposal, include the **proposal** statement and specify an IPsec proposal name at the **[edit services ipsec-vpn ipsec]** hierarchy level:

```

[edit services ipsec-vpn ipsec]
proposal proposal-name {
  authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
  description description;
  encryption-algorithm algorithm;
  lifetime-seconds seconds;
  protocol (ah | esp | bundle);
}

```

This section discusses the following topics:

- [Configuring the Authentication Algorithm for an IPsec Proposal on page 416](#)
- [Configuring the Description for an IPsec Proposal on page 418](#)
- [Configuring the Encryption Algorithm for an IPsec Proposal on page 418](#)
- [Configuring the Lifetime for an IPsec SA on page 418](#)
- [Configuring the Protocol for a Dynamic SA on page 419](#)

Configuring the Authentication Algorithm for an IPsec Proposal

To configure the authentication algorithm for an IPsec proposal, include the **authentication-algorithm** statement at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec proposal proposal-name]  
authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
```

The authentication algorithm can be one of the following:

- **hmac-md5-96**—Hash algorithm that authenticates packet data. It produces a 128-bit digest. Only 96 bits are used for authentication.
- **hmac-sha1-96**—Hash algorithm that authenticates packet data. It produces a 160-bit digest. Only 96 bits are used for authentication.



NOTE: Keep the following points in mind when you configure the authentication algorithm in an IPsec proposal:

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, an error occurs and the tunnel is not established in this scenario. For example, if one end of the tunnel contains router 1 configured with the authentication algorithm as hmac-sha-256-128 and the other end of the tunnel contains router 2 configured with the authentication algorithm as hmac-md5-96, the VPN tunnel is not established.
- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, and when one end of the tunnel contains two IPsec proposals to check whether a less secure algorithm is selected or not, an error occurs and the tunnel is not established. For example, if you configure two authentication algorithms for an IPsec proposal as hmac-sha-256-128 and hmac-md5-96 on one end of the tunnel, router 1, and if you configure the algorithm for an IPsec proposal as hmac-md5-96 on the other end of the tunnel, router 2, the tunnel is not established and the number of proposals mismatch.
- When you configure two IPsec proposals at both ends of a tunnel, such as the authentication-algorithm hmac-sha-256-128 and authentication-algorithm hmac-md5-96 statements at the [edit services ipsec-vpn ipsec proposal *proposal-name*] hierarchy level on one of the tunnel, router 1 (with the algorithms in two successive statements to specify the order), and the authentication-algorithm hmac-md5-96 and authentication-algorithm hmac-sha-256-128 statements at the [edit services ipsec-vpn ipsec proposal *proposal-name*] hierarchy level on one of the tunnel, router 2 (with the algorithms in two successive statements to specify the order, which is the reverse order of router 1), the tunnel is established in this combination as expected because the number of proposals is the same on both ends and they contain the same set of algorithms. However, the authentication algorithm selected is hmac-md5-96 and not the stronger algorithm of hmac-sha-256-128. This method of selection of the algorithm occurs because the first matching proposal is selected. Also, for a default proposal, regardless of whether the router supports the Advanced Encryption Standard (AES) encryption algorithm, the 3des-cbc algorithm is chosen and not the aes-cfb algorithm, which is because of the first algorithm in the default proposal being selected. In the sample scenario described here, on router 2, if you reverse the order of the algorithm configuration in the proposal so that it is the same order as the one specified on router 1, hmac-sha-256-128 is selected as the authentication method.
- You must be aware of the order of proposals in an IPsec policy at the time of configuration if you want the matching of proposals to happen in a certain order of preference, such as the strongest algorithm to be considered first when a match is made when both policies from the two peers have a proposal.

Configuring the Description for an IPsec Proposal

To specify an optional text description for an IPsec proposal, include the **description** statement at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec proposal proposal-name]  
description description;
```

Configuring the Encryption Algorithm for an IPsec Proposal

To configure encryption algorithm for an IPsec proposal, include the **encryption-algorithm** statement at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec proposal proposal-name]  
encryption-algorithm algorithm;
```

The encryption algorithm can be one of the following:

- **3des-cbc**—Encryption algorithm that has a block size of 24 bytes; its key size is 192 bits long.
- **des-cbc**—Encryption algorithm that has a block size of 8 bytes; its key size is 48 bits long.
- **aes-128-cbc**—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- **aes-192-cbc**—Advanced Encryption Standard (AES) 192-bit encryption algorithm.
- **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption algorithm.



NOTE: For a list of Data Encryption Standard (DES) encryption algorithm weak and semiweak keys, see RFC 2409, *The Internet Key Exchange (IKE)*. The AES encryption algorithms use a software implementation that has much lower throughput, so DES remains the recommended option.

For **3des-cbc**, the first 8 bytes should differ from the second 8 bytes, and the second 8 bytes should be the same as the third 8 bytes.

If you do not configure specific authentication or encryption settings, Junos OS uses the default values of **sha1** for the authentication and **3des-cbc** for the encryption. For NULL encryption to be effective, you must always specify the Encapsulating Security Payload (ESP) protocol for the NULL encryption algorithm by including the **protocol esp** statement at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** hierarchy level, regardless of other system configurations.

Configuring the Lifetime for an IPsec SA

When a dynamic IPsec SA is created, two types of lifetimes are used: hard and soft. The hard lifetime specifies the lifetime of the SA. The soft lifetime, which is derived from the hard lifetime, informs the IPsec key management system that the SA is about to expire.

This allows the key management system to negotiate a new SA before the hard lifetime expires.



NOTE: In IKEv1, the lifetime for SAs is negotiated with the remote peer based on the type of lifetime configured in the IPsec proposal. In IKEv2, such a negotiation is not performed with the remote peer. Instead, each IKE peer uses the lifetime that is locally configured for them.

For SAs in IKEv2, the lifetime is either the default value as IKEv1 (if another lifetime is not configured in the IPsec proposal) or all IKEv2 proposals in the IPsec policy must be configured with the same lifetime value.

To configure the hard lifetime value, include the **lifetime-seconds** statement and specify the number of seconds at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec proposal proposal-name]
lifetime-seconds seconds;
```

The default lifetime is 28,800 seconds. The range is from 180 through 86,400 seconds.

The soft lifetime values are as follows:

- Initiator: Soft lifetime = Hard lifetime – 135 seconds.
- Responder: Soft lifetime = Hard lifetime – 90 seconds.

Configuring the Protocol for a Dynamic SA

The **protocol** statement sets the protocol for a dynamic SA. IPsec uses two protocols to protect IP traffic: ESP and AH. The ESP protocol can support authentication, encryption, or both. The AH protocol is used for strong authentication. AH also authenticates the IP packet. The **bundle** option uses AH authentication and ESP encryption; it does not use ESP authentication because AH provides stronger authentication of IP packets.

To configure the protocol for a dynamic SA, include the **protocol** statement and specify the **ah**, **esp**, or **bundle** option at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec proposal proposal-name]
protocol (ah | esp | bundle);
```

Related Documentation

- [Configuring IPsec Policies on page 420](#)
- [Configuring IKE Proposals on page 405](#)
- [Configuring IKE Policies on page 409](#)
- [Configuring Security Associations on page 385](#)

Configuring IPsec Policies

An IPsec policy defines a combination of security parameters (IPsec proposals) used during IPsec negotiation. It defines Perfect Forward Secrecy (PFS) and the proposals needed for the connection. During the IPsec negotiation, IPsec looks for a proposal that is the same on both peers. The peer that initiates the negotiation sends all its policies to the remote peer, and the remote peer tries to find a match.

A match is made when both policies from the two peers have a proposal that contains the same configured attributes. If the lifetimes are not identical, the shorter lifetime between the two policies (from the host and peer) is used.

You can create multiple, prioritized IPsec proposals at each peer to ensure that at least one proposal matches a remote peer's proposal.

First, you configure one or more IPsec proposals; then you associate these proposals with an IPsec policy. You can prioritize a list of proposals used by IPsec in the **policy** statement by listing the proposals you want to use, from first to last.

To configure an IPsec policy, include the **policy** statement, and specify the policy name and one or more proposals to associate with the policy, at the **[edit services ipsec-vpn ipsec]** hierarchy level:

```
[edit services ipsec-vpn ipsec]
policy policy-name {
  description description;
  perfect-forward-secrecy {
    keys (group1 | group2 | group5 | group14);
  }
  proposals [ proposal-names ];
}
```

This section includes the following topics related to configuring an IPsec policy:

- [Configuring the Description for an IPsec Policy on page 420](#)
- [Configuring Perfect Forward Secrecy on page 421](#)
- [Configuring the Proposals in an IPsec Policy on page 421](#)
- [IPsec Policy for Dynamic Endpoints on page 421](#)
- [Example: Configuring an IPsec Policy on page 422](#)

Configuring the Description for an IPsec Policy

To specify an optional text description for an IPsec policy, include the **description** statement at the **[edit services ipsec-vpn ipsec policy policy-name]** hierarchy level:

```
[edit services ipsec-vpn ipsec policy policy-name]
description description;
```

Configuring Perfect Forward Secrecy

Perfect Forward Secrecy (PFS) provides additional security by means of a Diffie-Hellman shared secret value. With PFS, if one key is compromised, previous and subsequent keys are secure because they are not derived from previous keys. This statement is optional.

To configure PFS, include the **perfect-forward-secrecy** statement and specify a Diffie-Hellman group at the **[edit services ipsec-vpn ipsec policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec policy policy-name]
perfect-forward-secrecy {
  keys (group1 | group2 | group5 | group14);
}
```

The key can be one of the following:

- **group1**—Specifies that IKE use the 768-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- **group2**—Specifies that IKE use the 1024-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- **group5**—Specifies that IKE use the 1536-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- **group14**—Specifies that IKE use the 2048-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.

The higher numbered groups provide more security than the lowered numbered groups, but require more processing time.

Configuring the Proposals in an IPsec Policy

The IPsec policy includes a list of one or more proposals associated with an IPsec policy.

To configure the proposals in an IPsec policy, include the **proposals** statement and specify one or more proposal names at the **[edit services ipsec-vpn ipsec policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec policy policy-name]
proposals [ proposal-names ];
```

IPsec Policy for Dynamic Endpoints

An IPsec policy for dynamic endpoints defines a combination of security parameters (IPsec proposals) used during IPsec negotiation between dynamic peer security gateways, in which the remote ends of tunnels do not have a statically assigned IP address. During the IPsec negotiation, the IPsec policy looks for an IPsec proposal that is the same on both peers. The peer that initiates the negotiation sends all its policies to the remote peer, and the remote peer tries to find a match. A match is made when the policies from the two peers have a proposal that contains the same configured attributes. If the lifetimes are not identical, the shorter lifetime between the two policies (from the host and peer) is used.

If no policy is set, any policy proposed by the dynamic peer is accepted.

Example: Configuring an IPsec Policy

Define an IPsec policy, **dynamic policy-1**, that is associated with two proposals (**dynamic-1** and **dynamic-2**):

```
[edit services ipsec-vpn ipsec]
proposal dynamic-1 {
  protocol esp;
  authentication-algorithm hmac-md5-96;
  encryption-algorithm 3des-cbc;
  lifetime-seconds 6000;
}
proposal dynamic-2 {
  protocol esp;
  authentication-algorithm hmac-sha1-96;
  encryption-algorithm 3des-cbc;
  lifetime-seconds 6000;
}
policy dynamic-policy-1 {
  perfect-forward-secrecy {
    keys group1;
  }
  proposals [ dynamic-1 dynamic-2 ];
}
```



NOTE: Updates to the current IPsec proposal and policy configuration are not applied to the current IPsec SA; updates are applied to new IPsec SAs.

If you want the new updates to take immediate effect, you must clear the existing IPsec security associations so that they will be reestablished with the changed configuration. For information about how to clear the current IPsec security association, see the [Junos OS System Basics and Services Command Reference](#).

- Related Documentation**
- [Configuring IPsec Proposals on page 415](#)
 - [Configuring IKE Proposals on page 405](#)
 - [Configuring IKE Policies on page 409](#)
 - [Configuring Security Associations on page 385](#)

Configuring IPsec Rules

To configure an IPsec rule, include the **rule** statement and specify a rule name at the **[edit services ipsec-vpn]** hierarchy level:

```
[edit services ipsec-vpn]
rule rule-name {
  match-direction (input | output);
}
```



```

term term-name {
  from {
    destination-address address;
    ipsec-inside-interface interface-name;
    source-address address;
  }
  then {
    anti-replay-window-size bits;
    backup-remote-gateway address;
    clear-dont-fragment-bit;
    dynamic {
      ike-policy policy-name;
      ipsec-policy policy-name;
    }
    initiate-dead-peer-detection;
    dead-peer-detection {
      interval seconds;
      threshold number;
    }
    manual {
      direction (inbound | outbound | bidirectional) {
        authentication {
          algorithm (hmac-md5-96 | hmac-sha1-96);
          key (ascii-text key | hexadecimal key);
        }
        auxiliary-spi spi-value;
        encryption {
          algorithm algorithm;
          key (ascii-text key | hexadecimal key);
        }
        protocol (ah | bundle | esp);
        spi spi-value;
      }
    }
    no-anti-replay;
    remote-gateway address;
    syslog;
    tunnel-mtu bytes;
  }
}

```

Each IPsec rule consists of a set of terms, similar to a firewall filter.

A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded.

- **then** statement—Specifies the actions and action modifiers to be performed by the router software.

The following sections explain how to configure the components of IPsec rules:

- [Configuring Match Direction for IPsec Rules on page 424](#)
- [Configuring Match Conditions in IPsec Rules on page 424](#)
- [Configuring Actions in IPsec Rules on page 426](#)

Configuring Match Direction for IPsec Rules

Each rule must include a **match-direction** statement that specifies whether the match is applied on the input or output side of the interface. To configure where the match is applied, include the **match-direction (input | output)** statement at the **[edit services ipsec-vpn rule *rule-name*]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name]  
match-direction (input | output);
```

The match direction is used with respect to the traffic flow through the AS or Multiservices PIC. When a packet is sent to the PIC, direction information is carried along with it.

With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.

With a next-hop service set, packet direction is determined by the interface used to route the packet to the AS or Multiservices PIC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the PIC, the packet direction is output.

On the AS or Multiservices PIC, a flow lookup is performed. If no flow is found, rule processing is performed. All rules in the service set are considered. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that matches the packet direction are considered.

Configuring Match Conditions in IPsec Rules

To configure the match conditions in an IPsec rule, include the **from** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name*]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name]  
from {  
  destination-address address;  
  ipsec-inside-interface interface-name;  
  source-address address;  
}
```

You can use either the source address or the destination address as a match condition, in the same way that you would configure a firewall filter; for more information, see the *Junos OS Routing Protocols Library for Routing Devices*.

IPsec services support both IPv4 and IPv6 address formats. If you do not specifically configure either the source address or destination address, the default value **0.0.0.0/0**

(IPv4 ANY) is used. To use IPv6 ANY (**0::0/128**) as either the source or destination address, you must configure it explicitly.

For next-hop-style service sets only, the **ipsec-inside-interface** statement allows you to assign a logical interface to the tunnels established as a result of this match condition. The **inside-service-interface** statement that you can configure at the **[edit services service-set name next-hop-service]** hierarchy level allows you to specify .1 and .2 as inside and outside interfaces. However, you can configure multiple adaptive services logical interfaces with the **service-domain inside** statement and use one of them to configure the **ipsec-inside-interface** statement.

The Junos OS evaluates the criteria you configure in the **from** statement. If multiple link-type tunnels are configured within the same next-hop-style service set, the **ipsec-inside-interface** value enables the rule lookup module to distinguish a particular tunnel from other tunnels in case the source and destination addresses for all of them are **0.0.0.0/0** (ANY-ANY).



NOTE: When you configure the **ipsec-inside-interface** statement, interface-style service sets are not supported.

A special situation is provided by a term containing an “any-any” match condition (usually because the **from** statement is omitted). If there is an any-any match in a tunnel, a flow is not needed, because all flows within this tunnel use the same security association (SA) and packet selectors do not play a significant role. As a result, these tunnels will use packet-based IPsec. This strategy saves some flow resources on the PIC, which can be used for other tunnels that need a flow-based service.

The following configuration example shows an any-any tunnel configuration with no **from** statement in **term-1**. Missing selectors in the **from** clause result in a packet-based IPsec service.

```
services {
  ipsec-vpn {
    rule rule-1 {
      term term-1 {
        then {
          remote-gateway 10.1.0.1;
          dynamic {
            ike-policy ike_policy;
            ipsec-policy ipsec_policy;
          }
        }
      }
    }
    match-direction input;
  }
  .....
}
```

Flowless IPsec service is provided to link-type tunnels with an any-any matching, as well as to dynamic tunnels with any-any matching in both dedicated and shared mode.

For link-type tunnels, a mixture of flowless and flow-based IPsec is supported within a service set. If a service set includes some terms with any-any matching and some terms with selectors in the **from** clause, packet-based service is provided for the any-any tunnels and flow-based service is provided for the other tunnels with selectors.

For non link-type tunnels, if a service set contains both any-any terms and selector-based terms, flow-based service is provided to all the tunnels.

Configuring Actions in IPsec Rules

To configure actions in an IPsec rule, include the **then** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name*]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name]
then {
  anti-replay-window-size bits;
  backup-remote-gateway address;
  clear-dont-fragment-bit;
  dynamic {
    ike-policy policy-name;
    ipsec-policy policy-name;
  }
  initiate-dead-peer-detection;
  dead-peer-detection {
    interval seconds;
    threshold number;
  }
  manual {
    direction (inbound | outbound | bidirectional) {
      authentication {
        algorithm (hmac-md5-96 | hmac-sha1-96);
        key (ascii-text key | hexadecimal key);
      }
      auxiliary-spi spi-value;
      encryption {
        algorithm algorithm;
        key (ascii-text key | hexadecimal key);
      }
      protocol (ah | bundle | esp);
      spi spi-value;
    }
  }
  no-anti-replay;
  remote-gateway address;
  syslog;
  tunnel-mtu bytes;
}
```

The principal IPsec actions are to configure a dynamic or manual SA:

- You configure a dynamic SA by including the **dynamic** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level and referencing policies you have configured at the **[edit services ipsec-vpn ipsec]** and **[edit services ipsec-vpn ike]** hierarchy levels.

- You configure a manual SA by including the **manual** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

You can configure the following additional properties:

- [Enabling IPsec Packet Fragmentation on page 427](#)
- [Configuring Destination Addresses for Dead Peer Detection on page 427](#)
- [Configuring or Disabling IPsec Anti-Replay on page 428](#)
- [Enabling System Log Messages on page 429](#)
- [Specifying the MTU for IPsec Tunnels on page 429](#)

Enabling IPsec Packet Fragmentation

To enable fragmentation of IP version 4 (IPv4) packets in IPsec tunnels, include the **clear-dont-fragment-bit** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
  clear-dont-fragment-bit;
```

Setting the **clear-dont-fragment-bit** statement clears the Don't Fragment (DF) bit in the packet header, regardless of the packet size. If the packet size exceeds the tunnel maximum transmission unit (MTU) value, the packet is fragmented before encapsulation. For IPsec tunnels, the default MTU value is 1500 regardless of the interface MTU setting.

Configuring Destination Addresses for Dead Peer Detection

To specify the remote address to which the IPsec traffic is directed, include the **remote-gateway** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
  remote-gateway address;
```

To specify a backup remote address, include the **backup-remote-gateway** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
  backup-remote-gateway address;
```

These two statements support both IPv4 and IPv6 address formats.

Configuring the **backup-remote-gateway** statement enables the dead peer detection (DPD) protocol, which monitors the tunnel state and remote peer availability. When the primary tunnel defined by the **remote-gateway** statement is active, the backup tunnel is in standby mode. If the DPD protocol determines that the primary remote gateway address is no longer reachable, a new tunnel is established to the backup address.

If there is no incoming traffic from a peer during a defined interval of 10 seconds, the router detects a tunnel as inactive. A global timer polls all tunnels every 10 seconds and the Adaptive Services (AS) or Multiservices Physical Interface Card (PIC) sends a message listing any inactive tunnels. If a tunnel becomes inactive, the router takes the following steps to fail over to the backup address:

1. The adaptive services message triggers the DPD protocol to send a hello message to the peer.
2. If no acknowledgment is received, two retries are sent at 2-second intervals, and then the tunnel is declared dead.
3. Failover takes place if the tunnel is declared dead or there is an IPsec Phase 1 negotiation timeout. The primary tunnel is put in standby mode and the backup becomes active.
4. If the negotiation to the backup tunnel times out, the router switches back to the primary tunnel. If both peers are down, it tries the failover six times. It then stops failing over and reverts to the original configuration, with the primary tunnel active and the backup in standby mode.

You can also enable triggering of DPD hello messages without configuring a backup remote gateway by including the **initiate-dead-peer-detection** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
initiate-dead-peer-detection;
dead-peer-detection {
  interval seconds;
  threshold number;
}
```

In addition, for IKEv1 SAs you can set **interval** and **threshold** options under the **dead-peer-detection** statement when using the **initiate-dead-peer-detection** statement. These options are not applicable to IKEv2 SAs, which will use the default values. The interval is the amount of time that the peer waits for traffic from its destination peer before sending a DPD request packet, and the threshold is the maximum number of unsuccessful DPD requests to be sent before the peer is considered unavailable.

The monitoring behavior is the same as described for the **backup-remote-gateway** statement. This configuration enables the router to initiate DPD hellos when a backup IPsec gateway does not exist, and clean up the IKE and IPsec SAs in case the IKE peer is not reachable.

If the DPD protocol determines that the primary remote gateway address is no longer reachable, a new tunnel is established to the backup address. However, when you configure **initiate-dead-peer-detection** without a backup remote gateway address and the DPD protocol determines that the primary remote gateway address is no longer reachable, the tunnel is declared dead and IKE and IPsec SAs are cleaned up.

For more information on the DPD protocol, see RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*.

Configuring or Disabling IPsec Anti-Replay

To configure the size of the IPsec antireplay window, include the **anti-replay-window-size** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
```

`anti-replay-window-size bits;`

anti-replay-window-size can take values in the range from 64 through 4096 bits. The default value is 64 bits for AS PICs and 128 bits for Multiservices PICs and DPCs. AS PICs can support a maximum replay window size of 1024 bits, whereas Multiservices PICs and DPCs can support a maximum replay window size of 4096 bits. When the software is committing an IPsec configuration, the key management process (kmd) is unable to differentiate between the service interface types. As a result, if the maximum antireplay window size exceeds 1024 for AS PICs, the commit succeeds and no error message is produced. However, the software internally sets the antireplay window size for AS PICs to 1024 bits even if the configured value of the **anti-replay-window-size** is larger.

To disable the IPsec antireplay feature, include the **no-anti-replay** statement at the **[edit services ipsec-vpn rule rule-name term term-name then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
no-anti-replay;
```

By default, antireplay service is enabled. Occasionally this can cause interoperability issues with other vendors' equipment.

Enabling System Log Messages

To record an alert in the system logging facility, include the **syslog** statement at the **[edit services ipsec-vpn rule rule-name term term-name then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
syslog;
```

Specifying the MTU for IPsec Tunnels

To configure a specific maximum transmission unit (MTU) value for IPsec tunnels, include the **tunnel-mtu** statement at the **[edit services ipsec-vpn rule rule-name term term-name then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
tunnel-mtu bytes;
```



NOTE: The **tunnel-mtu** setting is the only place you need to configure an MTU value for IPsec tunnels. Inclusion of an **mtu** setting at the **[edit interfaces sp-fpc/pic/port unit logical-unit-number family inet]** hierarchy level is not supported.

Related Documentation

- [Configuring IPsec Rule Sets on page 429](#)
- [Configuring Security Associations on page 385](#)

Configuring IPsec Rule Sets

The **rule-set** statement defines a collection of IPsec rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by

including the **rule-set** statement at the **[edit services ipsec-vpn]** hierarchy level with a **rule** statement for each rule:

```
[edit services ipsec-vpn]
rule-set rule-set-name {
  rule rule-name;
}
```

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules match the packet, the packet is dropped by default.

- Related Documentation**
- [Configuring IPsec Rules on page 422](#)
 - [Configuring Security Associations on page 385](#)

Service Sets for IPsec Tunnels

The Adaptive Services PIC supports two types of service sets when you configure IPsec tunnels. Because they are used for different purposes, it is important to know the differences between these service set types.

- **Next-hop service set**—Supports multicast and multicast-style dynamic routing protocols (such as OSPF) over IPsec. Next-hop service sets allow you to use *inside* and *outside* logical interfaces on the Adaptive Services PIC to connect with multiple routing instances. They also allow the use of Network Address Translation (NAT) and stateful firewall capabilities. However, next-hop service sets do not monitor Routing Engine traffic by default and require configuration of multiple service sets to support traffic from multiple interfaces.
- **Interface service set**—Applied to a physical interface and similar to a stateless firewall filter. They are easy to configure, can support traffic from multiple interfaces, and can monitor Routing Engine traffic by default. However, they cannot support dynamic routing protocols or multicast traffic over the IPsec tunnel.

In general, we recommend that you use next-hop service sets because they support routing protocols and multicast over the IPsec tunnel, they are easier to understand, and the routing table makes forwarding decisions without administrative intervention.

- Related Documentation**
- [Understanding Junos VPN Site Secure on page 371](#)
 - [Configuring Junos VPN Site Secure or IPsec VPN on page 507](#)

Configuring IPsec Service Sets

IPsec service sets require additional specifications that you configure at the **[edit services service-set service-set-name ipsec-vpn-options]** hierarchy level:

```
[edit services service-set service-set-name ipsec-vpn-options]
anti-replay-window-size bits;
```



```

clear-dont-fragment-bit;
copy-dont-fragment-bit
set-dont-fragment-bit
ike-access-profile profile-name;
local-gateway address;
no-anti-replay;
passive-mode-tunneling;
trusted-ca [ ca-profile-names ];
tunnel-mtu bytes;

```

Configuration of these statements is described in the following sections:

- [Configuring the Local Gateway Address for IPsec Service Sets on page 431](#)
- [Configuring IKE Access Profiles for IPsec Service Sets on page 432](#)
- [Configuring Certification Authorities for IPsec Service Sets on page 433](#)
- [Configuring or Disabling Antireplay Service on page 433](#)
- [Clearing the Don't-Fragment Bit on page 434](#)
- [Configuring Passive-Mode Tunneling on page 435](#)
- [Configuring the Tunnel MTU Value on page 436](#)

Configuring the Local Gateway Address for IPsec Service Sets

If you configure an IPsec service set, you must also configure a local IPv4 or IPv6 address by including the **local-gateway** statement:

- If the Internet Key Exchange (IKE) gateway IP address is in **inet.0** (the default situation), you configure the following statement:

```
local-gateway address;
```

- If the IKE gateway IP address is in a VPN routing and forwarding (VRF) instance, you configure the following statement:

```
local-gateway address routing-instance instance-name;
```

You can configure all the link-type tunnels that share the same local gateway address in a single next-hop-style service set. The value you specify for the **inside-service-interface** statement at the **[edit services service-set service-set-name]** hierarchy level should match the **ipsec-inside-interface** value, which you configure at the **[edit services ipsec-vpn rule rule-name term term-name from]** hierarchy level. For more information about IPsec configuration, see *Configuring IPsec Rules*.



NOTE: To configure link-type tunnels, you can configure AMS logical interfaces as the IPsec internal interfaces by using the **ipsec-inside-interface interface-name** statement at the **[edit services ipsec-vpn rule rule-name term term-name from]** hierarchy level.

IKE Addresses in VRF Instances

You can configure Internet Key Exchange (IKE) gateway IP addresses that are present in a VPN routing and forwarding (VRF) instance as long as the peer is reachable through the VRF instance.

For next-hop service sets, the key management process (kmd) places the IKE packets in the routing instance that contains the **outside-service-interface** value you specify, as in this example:

```
routing-instances vrf-nxthop {
  instance-type vrf;
  interface sp-1/1/0.2;
  ...
}
services service-set service-set-1 {
  next-hop-service {
    inside-service-interface sp-1/1/0.1;
    outside-service-interface sp-1/1/0.2;
  }
  ...
}
```

For interface service sets, the **service-interface** statement determines the VRF, as in this example:

```
routing-instances vrf-intf {
  instance-type vrf;
  interface sp-1/1/0.3;
  interface ge-1/2/0.1; # interface on which service set is applied
  ...
}
services service-set service-set-2 {
  interface-service {
    service-interface sp-1/1/0.3;
  }
  ...
}
```

Configuring IKE Access Profiles for IPsec Service Sets

For dynamic endpoint tunneling only, you need to reference the IKE access profile configured at the **[edit access]** hierarchy level. To do this, include the **ike-access-profile** statement at the **[edit services service-set service-set-name ipsec-vpn-options]** hierarchy level:

```
[edit services service-set service-set-name ipsec-vpn-options]
ike-access-profile profile-name;
```

The **ike-access-profile** statement must reference the same name as the **profile** statement you configured for IKE access at the **[edit access]** hierarchy level. You can reference only one access profile in each service set. This profile is used to negotiate IKE and IPsec security associations with dynamic peers only.



NOTE: If you configure an IKE access profile in a service set, no other service set can share the same **local-gateway** address.

Also, you must configure a separate service set for each VRF. All interfaces referenced by the **ipsec-inside-interface** statement within a service set must belong to the same VRF.

Configuring Certification Authorities for IPsec Service Sets

You can specify one or more trusted certification authorities by including the **trusted-ca** statement:

```
trusted-ca [ ca-profile-names ];
```

When you configure public key infrastructure (PKI) digital certificates in the IPsec configuration, each service set can have its own set of trusted certification authorities. The names you specify for the **trusted-ca** statement must match profiles configured at the **[edit security pki]** hierarchy level; for more information, see the *Junos OS Administration Library for Routing Devices*. For more information about IPsec digital certificate configuration, see *Configuring IPsec Rules*.

Configuring or Disabling Antireplay Service

You can include the **anti-replay-window-size** statement at the **[edit services service-set service-set-name ipsec-vpn-options]** hierarchy level to specify the size of the antireplay window.

```
anti-replay-window-size bits;
```

This statement is useful for dynamic endpoint tunnels for which you cannot configure the **anti-replay-window-size** statement at the **[edit services ipsec-vpn rule rule-name term term-name then]** hierarchy level.

For static IPsec tunnels, this statement sets the antireplay window size for all the static tunnels within this service set. If a particular tunnel needs a specific value for antireplay window size, set the **anti-replay-window-size** statement at the **[edit services ipsec-vpn rule rule-name term term-name then]** hierarchy level. If antireplay check has to be disabled for a particular tunnel in this service set, set the **no-anti-replay** statement at the **[edit services ipsec-vpn rule rule-name term term-name then]** hierarchy level.



NOTE: The **anti-replay-window-size** and **no-anti-replay** settings at the **[edit services ipsec-vpn rule rule-name term term-name then]** hierarchy level override the settings specified at the **[edit services service-set service-set-name ipsec-vpn-options]** hierarchy level.

You can also include the **no-anti-replay** statement at the **[edit services service-set service-set-name ipsec-vpn-options]** hierarchy level to disable IPsec antireplay service. It occasionally causes interoperability issues for security associations.

no-anti-replay;

This statement is useful for dynamic endpoint tunnels for which you cannot configure the **no-anti-replay** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

For static IPsec tunnels, this statement disables the antireplay check for all the tunnels within this service set. If antireplay check has to be enabled for a particular tunnel, then set the **anti-replay-window-size** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.



NOTE: Setting the **anti-replay-window-size** and **no-anti-replay** statements at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level overrides the settings specified at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level.

Clearing the Don't-Fragment Bit

You can include the **clear-dont-fragment-bit** statement at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level to clear the Don't Fragment (DF) bit on all IP version 4 (IPv4) packets entering the IPsec tunnel. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation.

clear-dont-fragment-bit;

This statement is useful for dynamic endpoint tunnels, for which you cannot configure the **clear-dont-fragment-bit** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

For static IPsec tunnels, setting this statement clears the DF bit on packets entering all the static tunnels within this service set. If you want to clear the DF bit on packets entering a specific tunnel, set the **clear-dont-fragment-bit** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

In packets that are transmitted through dynamic endpoint IPsec tunnels, you can enable the value set in the Don't Fragment (DF) bit of the packet entering the tunnel to be copied only to the outer header of the IPsec packet and to not cause any modification to the DF bit in the inner header of the IPsec packet. If the packet size exceeds the tunnel maximum transmission unit (MTU) value, the packet is fragmented before encapsulation. For IPsec tunnels, the default MTU value is 1500 regardless of the interface MTU setting. To copy the DF bit value to only the outer header and not modify the inner header, use the **copy-dont-fragment-bit** statement at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level. You can also configure the DF bit to be set only in the outer IPv4 header of the IPsec packet and not be defined in the inner IPv4 header. To configure the DF bit in only the outer header of the IPsec packet and to leave the inner header unmodified, include the **set-dont-fragment-bit** statement at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level. These settings apply for dynamic endpoint tunnels and not for static tunnels, for which you need to include the **copy-dont-fragment-bit** and **set-dont-fragment-bit** statements at the **[edit services**

`ipsec-vpn rule rule-name term term-name then`] hierarchy level to clear the DF bit in the IPv4 packets that enter the static tunnel. These functionalities are supported on MX Series routers with MS-MICs and MS-MPCs.

Configuring Passive-Mode Tunneling

You can include the `passive-mode-tunneling` statement at the `[edit services service-set service-set-name ipsec-vpn-options]` hierarchy level to enable the service set to tunnel malformed packets.

```
[edit services service-set service-set-name ipsec-vpn-options]
passive-mode-tunneling;
```

This functionality bypasses the active IP checks, such as version, TTL, protocol, options, address and other land attack checks, and tunnels the packets as is. If this statement is not configured, packets failing the IP checks are dropped in the PIC. In passive mode, the inner packet is not touched; hence, an ICMP error is not generated, if the packet size exceeds the tunnel MTU value.

The IPsec tunnel is not treated as a next hop and TTL is not decremented. Because an ICMP error is not generated if the packet size exceeds the tunnel MTU value, the packet will be tunnelled even if it crosses the tunnel MTU threshold.



NOTE: This functionality is similar to that provided by the `no-ipsec-tunnel-in-traceroute` statement, described in *Disabling IPsec Tunnel Endpoint in Traceroute*. Starting with Junos OS Release 13.3R4 and 14.2R1, passive mode tunneling is supported on MS-MICs and MS-MPCs.



NOTE: The `header-integrity-check` option that is supported on MS-MICs and MS-MPCs to verify the packet header for anomalies in IP, TCP, UDP, and ICMP information and flag such anomalies and errors has a functionality that is opposite to the functionality caused by passive mode tunneling. If you configure both the `header-integrity-check` statement and the `passive-mode-tunneling` statement on MS-MICs and MS-MPCs, and attempt to commit such a configuration, an error is displayed during commit.

The passive mode tunneling functionality (by including the `passive-mode-tunneling` statement at the `[edit services service-set service-set-name ipsec-vpn-options]` hierarchy level) is a superset of the capability to disable IPsec tunnel endpoint in the traceroute output (by including `no-ipsec-tunnel-in-traceroute` statement at the `[edit services ipsec-vpn]` hierarchy level). Passive mode tunneling also bypasses the active IP checks and tunnel MTU check in addition to not treating an IPsec tunnel as a next-hop as configured by the `no-ipsec-tunnel-in-traceroute` statement.

Configuring the Tunnel MTU Value

You can include the **tunnel-mtu** statement at the **[edit services service-set service-set-name ipsec-vpn-options]** hierarchy level to set the maximum transmission unit (MTU) value for IPsec tunnels.

```
tunnel-mtu bytes;
```

This statement is useful for dynamic endpoint tunnels for which you cannot configure the **tunnel-mtu** statement at the **[edit services ipsec-vpn rule rule-name term term-name then]** hierarchy level.

For static IPsec tunnels, this statement sets the tunnel MTU value for all the tunnels within this service set. If you need a specific value for a particular tunnel, then set the **tunnel-mtu** statement at the **[edit services ipsec-vpn rule rule-name term term-name then]** hierarchy level.



NOTE: The **tunnel-mtu** setting at the **[edit services ipsec-vpn rule rule-name term term-name then]** hierarchy level overrides the value specified at the **[edit services service-set service-set-name ipsec-vpn-options]** hierarchy level.

Related Documentation

- [Understanding Service Sets on page 29](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 31](#)
- [Configuring Service Set Limitations on page 37](#)
- [Configuring System Logging for Service Sets on page 47](#)

Tracing Junos VPN Site Secure Operations



NOTE: Junos VPN Site Secure is a suite of IPsec features supported on multiservices line cards (MS-DPC, MS-MPC, and MS-MIC), and was previously referred to as IPsec services.

Trace operations track IPsec events and record them in a log file in the **/var/log** directory. By default, this file is named **/var/log/kmd**.

To trace IPsec operations, include the **traceoptions** statement at the **[edit services ipsec-vpn]** hierarchy level:

```
[edit services ipsec-vpn]
traceoptions {
  file <filename> <files number> <match regular-expression> <size bytes> <world-readable |
    no-world-readable>;
  flag flag;
  level level;
  no-remote-trace;
}
```

You can specify the following IPsec tracing flags:

- **all**—Trace everything.
- **certificates**—Trace certificates events.
- **database**—Trace security associations database events.
- **general**—Trace general events.
- **ike**—Trace IKE module processing.
- **parse**—Trace configuration processing.
- **policy-manager**—Trace policy manager processing.
- **routing-socket**—Trace routing socket messages.
- **snmp**—Trace SNMP operations.
- **timer**—Trace internal timer events.

The **level** statement sets the key management process (kmd) tracing level. The following values are supported:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match conditions that should be handled specially.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

This section includes the following topics:

- [Disabling IPsec Tunnel Endpoint in Traceroute on page 437](#)
- [Tracing IPsec PKI Operations on page 438](#)

Disabling IPsec Tunnel Endpoint in Traceroute

If you include the **no-ipsec-tunnel-in-traceroute** statement at the **[edit services ipsec-vpn]** hierarchy level, the IPsec tunnel is not treated as a next hop and the time to live (TTL) is not decremented. Also, if the TTL reaches zero, an ICMP time exceeded message is not generated.

```
[edit services ipsec-vpn]
no-ipsec-tunnel-in-traceroute;
```



NOTE: This functionality is also provided by the **passive-mode-tunneling** statement. You can use the **no-ipsec-tunnel-in-traceroute** statement in specific scenarios in which the IPsec tunnel should not be treated as a next hop and passive mode is not desired.

Tracing IPsec PKI Operations

Trace operations track IPsec PKI events and record them in a log file in the `/var/log` directory. By default, this file is named `/var/log/pkid`.

To trace IPsec PKI operations, include the **traceoptions** statement at the **[edit security pki]** hierarchy level:

```
[edit security pki]
traceoptions {
  file filename <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag (all | certificate-verification | enrollment | online-crl-check);
}
```

You can specify the following PKI tracing flags:

- **all**—Trace everything.
- **certificates**—Trace certificates events.
- **database**—Trace security associations database events.
- **general**—Trace general events.
- **ike**—Trace IKE module processing.
- **parse**—Trace configuration processing.
- **policy-manager**—Trace policy manager processing.
- **routing-socket**—Trace routing socket messages.
- **snmp**—Trace SNMP operations.
- **timer**—Trace internal timer events.

Related Documentation

- [Configuring IKE Policies on page 409](#)
- [Configuring IKE Proposals on page 405](#)

Multitask Example: Configuring IPsec Services

The following example-based instructions show how to configure IPsec services. The configuration involves defining an IKE policy, an IPsec policy, IPsec rules, trace options, and service sets.

This topic includes the following tasks:

1. [Configuring the IKE Proposal on page 439](#)
2. [Configuring the IKE Policy \(and Referencing the IKE Proposal\) on page 439](#)
3. [Configuring the IPsec Proposal on page 440](#)
4. [Configuring the IPsec Policy \(and Referencing the IPsec Proposal\) on page 441](#)
5. [Configuring the IPsec Rule \(and Referencing the IKE and IPsec Policies\) on page 441](#)

6. [Configuring IPsec Trace Options on page 442](#)
7. [Configuring the Access Profile \(and Referencing the IKE and IPsec Policies\) on page 443](#)
8. [Configuring the Service Set \(and Referencing the IKE Profile and the IPsec Rule\) on page 444](#)

Configuring the IKE Proposal

The IKE proposal configuration defines the algorithms and keys used to establish the secure IKE connection with the peer security gateway. For more information about IKE proposals, see *Configuring IKE Proposals*.

To define the IKE proposal:

1. In configuration mode, go to the following hierarchy level:

```
user@host# edit services ipsec-vpn
```
2. Configure the authentication method, which is **pre-shared keys** in this example:

```
[edit services ipsec-vpn]
user@host# set ike proposal test-IKE-proposal authentication-method pre-shared-keys
```
3. Configure the Diffie-Hellman Group and specify a name—for example, **group1**:

```
[edit services ipsec-vpn]
user@host# set ike proposal test-IKE-proposal dh-group group1
```
4. Configure the authentication algorithm, which is **sha1** in this example:

```
[edit services ipsec-vpn]
user@host# set ike proposal test-IKE-proposal authentication-algorithm sha1
```
5. Configure the encryption algorithm, which is **aes-256-cbc** in this example:

```
[edit services ipsec-vpn]
user@host# set ike proposal test-IKE-proposal encryption-algorithm aes-256-cbc
```

The following sample output shows the configuration of the IKE proposal:

```
[edit services ipsec-vpn]
user@host# show ike
proposal test-IKE-proposal {
    authentication-method pre-shared-keys;
    dh-group group1;
    authentication-algorithm sha1;
    encryption-algorithm aes-256-cbc;
}
```

Configuring the IKE Policy (and Referencing the IKE Proposal)

The IKE policy configuration defines the proposal, mode, addresses, and other security parameters used during IKE negotiation. For more information about IKE policies, see *Configuring IKE Policies*.

To define the IKE policy and reference the IKE proposal:

1. In configuration mode, go to the following hierarchy level:

```
user@host# edit services ipsec-vpn
```

2. Configure the IKE first phase mode—for example, **main**:

```
[edit services ipsec-vpn]
user@host# set ike policy test-IKE-policy mode main
```

3. Configure the proposal, which is **test-IKE-proposal** in this example:

```
[edit services ipsec-vpn]
user@host# set ike policy test-IKE-policy proposals test-IKE-proposal
```

4. Configure the local identification with an IPv4 address—for example, **192.168.255.2**:

```
[edit services ipsec-vpn]
user@host# set ike policy test-IKE-policy local-id ipv4_addr 192.168.255.2
```

5. Configure the preshared key in ASCII text format, which is **TEST** in this example:

```
[edit services ipsec-vpn]
user@host# set ike policy test-IKE-policy pre-shared-key ascii-text TEST
```

The following sample output shows the configuration of the IKE policy:

```
[edit services ipsec-vpn]
user@host# show ike
policy test-IKE-policy {
    mode main;
    proposals test-IKE-proposal;
    local-id ipv4_addr 192.168.255.2;
    pre-shared-key ascii-text TEST;
}
```

Configuring the IPsec Proposal

The IPsec proposal configuration defines the protocols and algorithms (security services) that are required to negotiate with the remote IPsec peer. For more information about IPsec proposals, see *Configuring IPsec Proposals*.

To define the IPsec proposal:

1. In configuration mode, go to the following hierarchy level:

```
user@host# edit services ipsec-vpn
```

2. Configure the IPsec protocol for the proposal—for example, **esp**:

```
[edit services ipsec-vpn]
user@host# set ipsec proposal test-IPsec-proposal protocol esp
```

3. Configure the authentication algorithm for the proposal, which is **hmac-sha1-96** in this example:

```
[edit services ipsec-vpn]
user@host# set ipsec proposal test-IPsec-proposal authentication-algorithm
hmac-sha1-96
```

4. Configure the encryption algorithm for the proposal, which is **aes-256-cbc** in this example:

```
[edit services ipsec-vpn]
```

```
user@host# set ipsec proposal test-IPsec-proposal encryption-algorithm aes-256-cbc
```

The following sample output shows the configuration of the IPsec proposal:

```
[edit services ipsec-vpn]
user@host# show ike
proposal test-IPsec-proposal {
  protocol esp;
  authentication-algorithm hmac-sha1-96;
  encryption-algorithm aes-256-cbc;
}
```

Configuring the IPsec Policy (and Referencing the IPsec Proposal)

The IPsec policy configuration defines a combination of security parameters (IPsec proposals) used during IPsec negotiation. It defines PFS and the proposals needed for the connection. For more information about IPsec policies, see *Configuring IPsec Policies*.

To define the IPsec policy and reference the IPsec proposal:

1. In configuration mode, go to the following hierarchy level:

```
user@host# edit services ipsec-vpn
```

2. Configure the keys for perfect forward secrecy in the IPsec policy—for example, **group1**:

```
[edit services ipsec-vpn]
user@host# set ipsec policy test-IPsec-policy perfect-forward-secrecy keys group1
```

3. Configure a set of IPsec proposals in the IPsec policy—for example, **test-IPsec-proposal**:

```
[edit services ipsec-vpn]
user@host# set ipsec policy test-IPsec-policy proposals test-IPsec-proposal
```

The following sample output shows the configuration of the IPsec policy:

```
[edit services ipsec-vpn]
user@host# show ipsec policy test-IPsec-policy
perfect-forward-secrecy {
  keys group1;
}
proposals test-IPsec-proposal;
```

Configuring the IPsec Rule (and Referencing the IKE and IPsec Policies)

The IPsec rule configuration defines the direction that specifies whether the match is applied on the input or output side of the interface. The configuration also consists of a set of terms that specify the match conditions and applications that are included and excluded and also specify the actions and action modifiers to be performed by the router software. For more information about IPsec rules, see *Configuring IPsec Rules*.

To define the IPsec rule and reference the IKE and IPsec policies:

1. In configuration mode, go to the following hierarchy level:

```
user@host# edit services ipsec-vpn
```

2. Configure the IP destination address for the IPsec term in the IPsec rule—for example, **192.168.255.2/32**:

```
[edit services ipsec-vpn]
user@host# set rule test-IPsec-rule term 10 from destination-address 192.168.255.2/32
```

3. Configure the remote gateway address for the IPsec term in the IPsec rule—for example, **0.0.0.0**:

```
[edit services ipsec-vpn]
user@host# set rule test-IPsec-rule term 10 then remote-gateway 0.0.0.0
```

4. Configure a dynamic security association for IKE policy for the IPsec term in the IPsec rule, which is **test-IKE-policy** in this example:

```
[edit services ipsec-vpn]
user@host# set rule test-IPsec-rule term 10 then dynamic ike-policy test-IKE-policy
```

5. Configure a dynamic security association for IKE proposal for the IPsec term in the IPsec rule, which is **test-IPsec-proposal** in this example:

```
[edit services ipsec-vpn]
user@host# set rule test-IPsec-rule term 10 then dynamic ipsec-policy test-IPsec-policy
```

6. Configure a direction for which the rule match is being applied in the IPsec rule—for example, **input**:

```
[edit services ipsec-vpn]
user@host# set rule test-IPsec-rule match-direction input
```

The following sample output shows the configuration of the IPsec rule:

```
[edit services ipsec-vpn]
user@host# show rule test-IPsec-rule
term 10 {
  from {
    destination-address {
      192.168.255.2/32;
    }
  }
  then {
    remote-gateway 0.0.0.0;
    dynamic {
      ike-policy test-IKE-policy;
      ipsec-policy test-IPsec-policy;
    }
  }
}
match-direction input;
```

Configuring IPsec Trace Options

The IPsec trace options configuration tracks IPsec events and records them in a log file in the **/var/log** directory. By default, this file is named **/var/log/kmd**. For more information about IPsec rules, see *Tracing IPsec Operations*.

To define the IPsec trace options:

1. In configuration mode, go to the following hierarchy level:

```
user@host# edit services ipsec-vpn
```

2. Configure the trace file, which is **ipsec.log** in this example:

```
[edit services ipsec-vpn]
user@host# set traceoptions file ipsec.log
```

3. Configure all the tracing parameters with the option **all** in this example:

```
[edit services ipsec-vpn]
user@host# set traceoptions flag all
```

The following sample output shows the configuration of the IPsec trace options:

```
[edit services ipsec-vpn]
user@host# show traceoptions
file ipsec.log;
flag all;
```

Configuring the Access Profile (and Referencing the IKE and IPsec Policies)

The access profile configuration defines the access profile and references the IKE and IPsec policies. For more information about access profile, see *Configuring an IKE Access Profile*.

To define the access profile and reference the IKE and IPsec policies:

1. In configuration mode, go to the following hierarchy level:

```
user@host# [edit access]
```

2. Configure the list of local and remote proxy identity pairs with the **allowed-proxy-pair** option. In this example, **10.0.0.0/24** is the IP address for local proxy identity and **10.0.1.0/24** is the IP address for remote proxy identity:

```
[edit access]
user@host# set profile IKE-profile-TEST client * ike allowed-proxy-pair local
10.0.0.0/24 remote 10.0.1.0/24
```

3. Configure the IKE policy—for example, **test-IKE-policy**:

```
[edit access]
user@host# set profile IKE-profile-TEST client * ike ike-policy test-IKE-policy
```

4. Configure the IPsec policy—for example, **test-IPsec-policy**:

```
[edit access]
user@host# set profile IKE-profile-TEST client * ike ipsec-policy test-IPsec-policy
```

5. Configure the identity of logical service interface pool, which is **TEST-intf** in this example:

```
[edit access]
user@host# set profile IKE-profile-TEST client * ike interface-id TEST-intf
```

The following sample output shows the configuration of the access profile:

```
[edit access]
user@host# show
profile IKE-profile-TEST {
```

```
client * {  
    ike {  
        allowed-proxy-pair local 10.0.0.0/24 remote 10.0.1.0/24;  
        ike-policy test-IKE-policy;  
        ipsec-policy test-IPsec-policy; # new statement  
        interface-id TEST-intf;  
    }  
}
```

Configuring the Service Set (and Referencing the IKE Profile and the IPsec Rule)

The service set configuration defines IPsec service sets that require additional specifications and references the IKE profile and the IPsec rule. For more information about IPsec service sets, see [“Configuring IPsec Service Sets” on page 430](#).

To define the service set configuration with the next-hop service sets and IPsec VPN options:

1. In configuration mode, go to the following hierarchy level:

```
user@host# [edit services]
```

2. Configure a service set with parameters for next hop service interfaces for the inside network—for example, **sp-1/2/0.1**:

```
[edit services]
```

```
user@host# set service-set TEST next-hop-service inside-service-interface sp-1/2/0.1
```

3. Configure a service set with parameters for next hop service interfaces for the outside network—for example, **sp-1/2/0.2**:

```
[edit services]
```

```
user@host# set service-set TEST next-hop-service outside-service-interface sp-1/2/0.2
```

4. Configure the IPsec VPN options with the address and routing instance for the local gateway—for example, **192.168.255.2**:

```
[edit services]
```

```
user@host# set service-set TEST ipsec-vpn-options local-gateway 192.168.255.2
```

5. Configure the IPsec VPN options with the IKE access profile for dynamic peers, which is **IKE-profile-TEST** in this example:

```
[edit services]
```

```
user@host# set service-set TEST ipsec-vpn-options ike-access-profile IKE-profile-TEST
```

6. Configure a service set with IPsec VPN rules, which is **test-IPsec-rule** in this example:

```
[edit services]
```

```
user@host# set service-set TEST ipsec-vpn-rules test-IPsec-rule
```

The following sample output shows the configuration of the service set configuration referencing the IKE profile and the IPsec rule:

```
[edit services]user@host# show service-set TEST  
next-hop-service {  
    inside-service-interface sp-1/2/0.1;  
    outside-service-interface sp-1/2/0.2;  
}
```

```
ipsec-vpn-options {  
    local-gateway 192.168.255.2;  
    ike-access-profile IKE-profile-TEST;  
}  
ipsec-vpn-rules test-IPsec-rule;
```

- Related Documentation**
- *Configuring IKE Proposals*
 - *Configuring IKE Policies*
 - *Configuring IPsec Proposals*
 - *Configuring IPsec Policies*
 - *Configuring IPsec Rules*
 - *Tracing IPsec Operations*
 - *Configuring an IKE Access Profile*
 - [Configuring IPsec Service Sets on page 430](#)

CHAPTER 32

Enhancing Security with Static IPsec over VRF

- [Example: Configuring Statically Assigned IPsec Tunnels over a VRF Instance on page 447](#)

Example: Configuring Statically Assigned IPsec Tunnels over a VRF Instance

This example shows how to configure a statically assigned IPsec tunnel over a VRF instance, and contains the following sections:

- [Requirements on page 447](#)
- [Overview on page 447](#)
- [Configuration on page 447](#)
- [Verification on page 453](#)

Requirements

This example uses the following hardware and software components:

- M Series, MX Series, or T Series router that is configured as a provider edge router.
- Junos OS Release 9.4 and later.

No special configuration beyond device initialization is required before you can configure this feature.

Overview

Junos OS enables you to configure statically assigned IPsec tunnels on Virtual Routing and Forwarding (VRF) instances. Ability to configure IPsec tunnels on VRF instances enhances network segmentation and security. You can have multiple customer tunnels configured on the same PE router over VRF instances. Each VRF instance acts as logical router with an exclusive routing table.

Configuration

This example shows the configuration of an IPsec tunnel over a VRF instance on a provider edge router, and provides step-by-step instructions for completing the required configuration.

This section contains:

- [Configuring the Provider Edge Router on page 448](#)
- [Results on page 450](#)

Configuring the Provider Edge Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

```
set interfaces ge-0/3/0 unit 0 family inet address 10.6.6.6/32
set interfaces ge-1/1/0 description "teller ge-0/1/0"
set interfaces ge-1/1/0 unit 0 family inet address 10.21.1.1/16
set interfaces ms-1/2/0 unit 0 family inet address 10.7.7.7/32
set interfaces ms-1/2/0 unit 1 family inet
set interfaces ms-1/2/0 unit 1 service-domain inside
set interfaces ms-1/2/0 unit 2 family inet
set interfaces ms-1/2/0 unit 2 service-domain outside
set policy-options policy-statement vpn-export then community add vpn-community
set policy-options policy-statement vpn-export then accept
set policy-options policy-statement vpn-import term a from community vpn-community
set policy-options policy-statement vpn-import term a then accept
set policy-options community vpn-community members target:100:20
set routing-instances vrf instance-type vrf
set routing-instances vrf interface ge-0/3/0.0
set routing-instances vrf interface ms-1/2/0.1
set routing-instances vrf route-distinguisher 192.168.0.1:1
set routing-instances vrf vrf-import vpn-import
set routing-instances vrf vrf-export vpn-export
set routing-instances vrf routing-options static route 10.0.0.0/0 next-hop ge-0/3/0.0
set routing-instances vrf routing-options static route 10.11.11.1/32 next-hop ge-0/3/0.0
set routing-instances vrf routing-options static route 10.8.8.1/32 next-hop ms-1/2/0.1
set services ipsec-vpn ipsec proposal demo_ipsec_proposal protocol esp
set services ipsec-vpn ipsec proposal demo_ipsec_proposal authentication-algorithm
    hmac-sha1-96
set services ipsec-vpn ipsec proposal demo_ipsec_proposal encryption-algorithm 3des-cbc
set services ipsec-vpn ipsec policy demo_ipsec_policy perfect-forward-secrecy keys
    group2
set services ipsec-vpn ipsec policy demo_ipsec_policy proposals demo_ipsec_proposal
set services ipsec-vpn ike proposal demo_ike_proposal authentication-method
    pre-shared-keys
set services ipsec-vpn ike proposal demo_ike_proposal dh-group group2
set services ipsec-vpn ike policy demo_ike_policy proposals demo_ike_proposal
set services ipsec-vpn ike policy demo_ike_policy pre-shared-key ascii-text juniperkey
set services ipsec-vpn rule demo-rule term demo-term then remote-gateway 10.21.2.1
set services ipsec-vpn rule demo-rule term demo-term then dynamic ike-policy
    demo_ike_policy
set services ipsec-vpn rule demo-rule match-direction input
set services service-set demo-service-set next-hop-service inside-service-interface
    ms-1/2/0.1
set services service-set demo-service-set next-hop-service outside-service-interface
    ms-1/2/0.2
set services service-set demo-service-set ipsec-vpn-options local-gateway 10.21.1.1
```

```
set services service-set demo-service-set ipsec-vpn-rules demo-rule
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a statically assigned IPsec tunnel on a VRF instance:

1. Configure the interfaces. In this step, you configure two Ethernet (**ge**) interfaces, one services interface (**ms**-), and also the service-domain properties for the logical interfaces of the services interface. Note that the logical interface that is marked as the inside interface applies the configured service on the traffic, whereas the one that is marked as the outside interface acts as the egress point for the traffic on which the inside interface has applied the service.

```
[edit interfaces]
user@PE1# set ge-0/3/0 unit 0 family inet address 10.6.6/32
user@PE1# set ge-1/1/0 description "teller ge-0/1/0"
user@PE1# set ge-1/1/0 unit 0 family inet address 10.21.1/16
user@PE1# set ms-1/2/0 unit 0 family inet address 10.7.7/32
user@PE1# set ms-1/2/0 unit 1 family inet
user@PE1# set ms-1/2/0 unit 1 service-domain inside
user@PE1# set ms-1/2/0 unit 2 family inet
user@PE1# set ms-1/2/0 unit 2 service-domain outside
```

2. Configure a routing policy to specify route import and export criteria for the VRF instance. The import and export policies defined in this step are referenced from the routing-instance configuration in the next step.

```
[edit policy-options]
user@PE1# set policy-statement vpn-export then community add vpn-community
user@PE1# set policy-statement vpn-export then accept
user@PE1# set policy-statement vpn-import term a from community vpn-community
user@PE1# set policy-statement vpn-import term a then accept
user@PE1# set community vpn-community members target:100:20
```

3. Configure a routing instance and specify the routing-instance type as **vrf**. Apply the import and export policies defined in the previous step to the routing instance, and specify a static route to send the IPsec traffic to the inside interface (**ms-1/2/0.1**) configured in the first step.

```
[edit routing-instance]
user@PE1# set vrf instance-type vrf
user@PE1# set vrf interface ge-0/3/0.0
user@PE1# set vrf interface ms-1/2/0.1
user@PE1# set vrf route-distinguisher 192.168.0.1:1
user@PE1# set vrf vrf-import vpn-import
user@PE1# set vrf vrf-export vpn-export
user@PE1# set vrf routing-options static route 10.0.0.0/0 next-hop ge-0/3/0.0
user@PE1# set vrf routing-options static route 10.11.1/32 next-hop ge-0/3/0.0
user@PE1# set vrf routing-options static route 10.8.8.1/32 next-hop ms-1/2/0.1
```

4. Configure IKE and IPsec proposals and policies, and a rule to apply the IKE policy on the incoming traffic..



NOTE: By default, Junos OS uses IKE policy version 1.0. Junos OS Release 11.4 and later also support IKE policy version 2.0 which you must configure at `[edit services ipsec-vpn ike policy policy-name pre-shared]`.

```
[edit services]
user@PE1# set ipsec-vpn ipsec proposal demo_ipsec_proposal protocol esp
user@PE1# set ipsec-vpn ipsec proposal demo_ipsec_proposal
  authentication-algorithm hmac-sha1-96
user@PE1# set ipsec-vpn ipsec proposal demo_ipsec_proposal encryption-algorithm
  3des-cbc
user@PE1# set ipsec-vpn ipsec policy demo_ipsec_policy perfect-forward-secrecy
  keys group2
user@PE1# set ipsec-vpn ipsec policy demo_ipsec_policy proposals
  demo_ipsec_proposal
user@PE1# set ipsec-vpn ike proposal demo_ike_proposal authentication-method
  pre-shared-keys
user@PE1# set ipsec-vpn ike proposal demo_ike_proposal dh-group group2
user@PE1# set ipsec-vpn ike policy demo_ike_policy proposals demo_ike_proposal
user@PE1# set ipsec-vpn ike policy demo_ike_policy pre-shared-key ascii-text
  juniperkey
user@PE1# set ipsec-vpn rule demo-rule term demo-term then remote-gateway
  10.21.2.1
user@PE1# set ipsec-vpn rule demo-rule term demo-term then dynamic ike-policy
  demo_ike_policy
user@PE1# set ipsec-vpn rule demo-rule match-direction input
```

5. Configure a next-hop style service set. Note that you must configure the inside and outside interfaces that you configured in the first step as the **inside-service-interface** and **outside-service-interface** respectively.

```
[edit services]
user@PE1# set service-set demo-service-set next-hop-service
  inside-service-interface ms-1/2/0.1
user@PE1# set service-set demo-service-set next-hop-service
  outside-service-interface ms-1/2/0.2
user@PE1# set service-set demo-service-set ipsec-vpn-options local-gateway
  10.21.1.1
user@PE1# set service-set demo-service-set ipsec-vpn-rules demo-rule
```

6. Commit the configuration.

```
[edit]
user@PE1# commit
```

Results

From the configuration mode of Router 1, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show routing-instances**, **show services ipsec-vpn**, and **show services service-set** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show interfaces
...
```

```

ms-1/2/0 {
  unit 0 {
    family inet {
      address 10.7.7/32;
    }
  }
  unit 1 {
    family inet;
    service-domain inside;
  }
  unit 2 {
    family inet;
    service-domain outside;
  }
}
ge-0/3/0 {
  unit 0 {
    family inet {
      address 10.6.6.6/32;
    }
  }
}
ge-1/1/0 {
  description "teller ge-0/1/0";
  unit 0 {
    family inet {
      address 10.21.1.1/16;
    }
  }
}
...
user@PE1# show policy-options
policy-statement vpn-export {
  then {
    community add vpn-community;
    accept;
  }
}
policy-statement vpn-import {
  term a {
    from community vpn-community;
    then accept;
  }
}
community vpn-community members target:100:20;

user@PE1# show routing-instances
vrf {
  instance-type vrf;
  interface ge-0/3/0.0;
  interface ms-1/2/0.1;
  route-distinguisher 192.168.0.1:1;
  vrf-import vpn-import;
  vrf-export vpn-export;
  routing-options {
    static {

```

```
        route 10.0.0.0/0 next-hop ge-0/3/0.0;
        route 10.11.11.1/32 next-hop ge-0/3/0.0;
        route 10.8.8.1/32 next-hop ms-1/2/0.1;
    }
}
}

user@PE1# show services ipsec-vpn
ipsec-vpn {
    rule demo-rule {
        term demo-term {
            then {
                remote-gateway 10.21.2.1;
                dynamic {
                    ike-policy demo_ike_policy;
                }
            }
        }
        match-direction input;
    }
    ipsec {
        proposal demo_ipsec_proposal {
            protocol esp;
            authentication-algorithm hmac-sha1-96;
            encryption-algorithm 3des-cbc;
        }
        policy demo_ipsec_policy {
            perfect-forward-secrecy {
                keys group2;
            }
            proposals demo_ipsec_proposal;
        }
    }
    ike {
        proposal demo_ike_proposal {
            authentication-method pre-shared-keys;
            dh-group group2;
        }
        policy demo_ike_policy {
            proposals demo_ike_proposal;
            pre-shared-key ascii-text "$9$JoUi.QF/0BEP5BEcyW8ZUjqPQ/9p0Ic"; ##
                SECRET-DATA
        }
    }
}

user@PE1# show services service-set demo-service-set
next-hop-service {
    inside-service-interface ms-1/2/0.1;
    outside-service-interface ms-1/2/0.2;
}
ipsec-vpn-options {
    local-gateway 10.21.1.1;
}
ipsec-vpn-rules demo-rule;
```

Verification

- [Verifying that the VRF instance is working on page 453](#)

Verifying that the VRF instance is working

Purpose**Action****Meaning**

- | | |
|----------------------|---|
| Related | • Understanding Junos VPN Site Secure on page 371 |
| Documentation | • Configuring Security Associations on page 385 |
| | • Configuring IPsec Proposals on page 415 |
| | • Configuring IKE Proposals on page 405 |

Dynamically Assigning Tunnels Using Junos VPN Site Secure

- [Configuring Dynamic Endpoints for IPsec Tunnels on page 455](#)
- [Example: Configuring Dynamically Assigned Policy Based Tunnels on page 461](#)
- [Example: Configuring IKE Dynamic SAs on page 466](#)
- [Example: IKE Dynamic SA Configuration with Digital Certificates on page 483](#)

Configuring Dynamic Endpoints for IPsec Tunnels

IPsec tunnels can also be established using *dynamic peer* security gateways, in which the remote ends of tunnels do not have a statically assigned IP address. Since the remote address is not known and might be pulled from an address pool each time the remote host reboots, establishment of the tunnel relies on using IKE **main** mode with either preshared global keys or digital certificates that accept any remote identification value. Both policy-based and link-type tunnels are supported:

- Policy-based tunnels used shared mode.
- Link-type or routed tunnels use dedicated mode. Each tunnel allocates a service interface from a pool of interfaces configured for the dynamic peers. Routing protocols can be configured to run on these service interfaces to learn routes over the IPsec tunnel that is used as a link in this scenario.

This section includes the following topics:

- [Authentication Process on page 456](#)
- [Implicit Dynamic Rules on page 456](#)
- [Reverse Route Insertion on page 457](#)
- [Configuring an IKE Access Profile on page 457](#)
- [Referencing the IKE Access Profile in a Service Set on page 459](#)
- [Configuring the Interface Identifier on page 459](#)
- [Default IKE and IPsec Proposals on page 460](#)

Authentication Process

The remote (dynamic peer) initiates the negotiations with the local (Juniper Networks) router. The local router uses the default IKE and IPsec policies to match the proposals sent by the remote peer to negotiate the security association (SA) values. Implicit proposals contain a list of all the supported transforms that the local router expects from all the dynamic peers.

If preshared key authentication is used, the preshared key is global for a service set. When seeking the preshared key for the peer, the local router matches the peer's source address against any explicitly configured preshared keys in that service set. If a match is not found, the local router uses the global preshared key for authentication. This key is the one configured in the IKE access profile referenced by the service set.

Phase 2 of the authentication matches the *proxy identities* of the protected hosts and networks sent by the peer against a list of configured proxy identities. The accepted proxy identity is used to create the dynamic rules for encrypting the traffic. You can configure proxy identities by including the **allowed-proxy-pair** statement in the IKE access profile. If no entry matches, the negotiation is rejected.

If you do not configure the **allowed-proxy-pair** statement, the default value **ANY(0.0.0.0/0)-ANY** is applied, and the local router accepts any proxy identities sent by the peer. Both IPv4 and IPv6 addresses are accepted, but you must configure all IPv6 addresses manually.

Once the phase 2 negotiation completes successfully, the router builds the dynamic rules and inserts the reverse route into the routing table using the accepted proxy identity.

Implicit Dynamic Rules

After successful negotiation with the dynamic peer, the key management process (kmd) creates a dynamic rule for the accepted phase 2 proxy and applies it on the local AS or Multiservices PIC. The source and destination addresses are specified by the accepted proxy. This rule is used to encrypt traffic directed to one of the end hosts in the phase 2 proxy identity.

The dynamic rule includes an **ipsec-inside-interface** value, which is the interface name assigned to the dynamic tunnel. The **source-address** and **destination-address** values are accepted from the proxy ID. The **match-direction** value is **input** for next-hop-style service sets.



NOTE: You do not configure this rule; it is created by the key management process (kmd).

Rule lookup for static tunnels is unaffected by the presence of a dynamic rule; it is performed in the order configured. When a packet is received for a service set, static rules are always matched first.

Dynamic rules are matched after the rule match for static rules has failed.

Response to dead peer detection (DPD) hello messages takes place the same way with dynamic peers as with static peers. Initiating DPD hello messages from dynamic peers is not supported.

Reverse Route Insertion

Static routes are automatically inserted into the route table for those networks and hosts protected by a remote tunnel endpoint. These protected hosts and networks are known as remote proxy identities.

Each route is created based on the remote proxy network and mask sent by the peer and is inserted in the relevant route table after successful phase 1 and phase 2 negotiations.

The route preference for each static reverse route is 1. This value is necessary to avoid conflict with similar routes that might be added by the routing protocol process (rpd).

No routes are added if the accepted remote proxy address is the default (0.0.0.0/0). In this case you can run routing protocols over the IPsec tunnel to learn routes and add static routes for the traffic you want to be protected over this tunnel.

For next-hop-style service sets, the reverse routes include next hops pointing to the locations specified by the **inside-service-interface** statement.

The route table in which to insert these routes depends on where the **inside-service-interface** location is listed. If these interfaces are present in a VPN routing and forwarding (VRF) instance, then routes are added to the corresponding VRF table; otherwise, the routes are added to **inet.0**.



NOTE: Reverse route insertion takes place only for tunnels to dynamic peers. These routes are added only for next-hop-style service sets.

Configuring an IKE Access Profile

You can configure only one tunnel profile per service set for all dynamic peers. The configured preshared key in the profile is used for IKE authentication of all dynamic peers terminating in that service set. Alternatively, you can include the **ike-policy** statement to reference an IKE policy you define with either specific identification values or a wildcard (the **any-remote-id** option). You configure the IKE policy at the **[edit services ipsec-vpn ike]** hierarchy level.

The IKE tunnel profile specifies all the information needed to complete the IKE negotiation. Each protocol has its own statement hierarchy within the client statement to configure protocol-specific attribute value pairs, but only one client configuration is allowed for each profile. The following is the configuration at the **[edit access]** hierarchy level; for more information on access profiles, see the *Junos OS Administration Library for Routing Devices*.

```
[edit access]
profile profile-name {
  client * {
    ike {
```

```

    allowed-proxy-pair {
        remote remote-proxy-address local local-proxy-address;
    }
    pre-shared-key (ascii-text key-string | hexadecimal key-string);
    ike-policy policy-name;
    interface-id <string-value>;
    ipsec-policy ipsec-policy;
}
}
}

```



NOTE: For dynamic peers, the Junos OS supports the IKE main mode with either the preshared key method of authentication or an IKE access profile that uses a local digital certificate.

- In preshared key mode, the IP address is used to identify a tunnel peer to get the preshared key information. The client value * (wildcard) means that configuration within this profile is valid for all dynamic peers terminating within the service set accessing this profile.
- In digital certificate mode, the IKE policy defines which remote identification values are allowed.

The following statements make up the IKE profile:

- **allowed-proxy-pair**—During phase 2 IKE negotiation, the remote peer supplies its network address (**remote**) and its peer's network address (**local**). Since multiple dynamic tunnels are authenticated through the same mechanism, this statement must include the list of possible combinations. If the dynamic peer does not present a valid combination, the phase 2 IKE negotiation fails.

By default, **remote 0.0.0.0/0 local 0.0.0.0/0** is used if no values are configured. Both IPv4 and IPv6 address formats are supported in this configuration, but there are no default IPv6 addresses. You must specify even **0::0/0**.

- **pre-shared-key**—Key used to authenticate the dynamic peer during IKE phase 1 negotiation. This key is known to both ends through an out-of-band secure mechanism. You can configure the value either in **hexadecimal** or **ascii-text** format. It is a mandatory value.
- **ike-policy**—Policy that defines the remote identification values corresponding to the allowed dynamic peers; can contain a wildcard value **any-remote-id** for use in dynamic endpoint configurations only.
- **interface-id**—Interface identifier, a mandatory attribute used to derive the logical service interface information for the session.
- **ipsec-policy**—Name of the IPsec policy that defines the IPsec policy information for the session. You define the IPsec policy at the **[edit services ipsec-vpn ipsec policy *policy-name*]** hierarchy level. If no policy is set, any policy proposed by the dynamic peer is accepted.

Referencing the IKE Access Profile in a Service Set

To complete the configuration, you need to reference the IKE access profile configured at the **[edit access]** hierarchy level. To do this, include the **ike-access-profile** statement at the **[edit services service-set name ipsec-vpn-options]** hierarchy level:

```
[edit services service-set name]
ipsec-vpn-options {
  local-gateway address;
  ike-access-profile profile-name;
}
next-hop-service {
  inside-service-interface interface-name;
  outside-service-interface interface-name;
}
```

The **ike-access-profile** statement must reference the same name as the **profile** statement you configured for IKE access at the **[edit access]** hierarchy level. You can reference only one access profile in each service set. This profile is used to negotiate IKE and IPsec security associations with dynamic peers only.



NOTE: If you configure an IKE access profile in a service set, no other service set can share the same local-gateway address.

Also, you must configure a separate service set for each VRF instance. All interfaces referenced by the **ipsec-inside-interface** statement within a service set must belong to the same VRF instance.

Configuring the Interface Identifier

You can configure an interface identifier for a group of dynamic peers, which specifies which adaptive services logical interface(s) take part in the dynamic IPsec negotiation. By assigning the same interface identifier to multiple logical interfaces, you can create a pool of interfaces for this purpose. To configure an interface identifier, include the **ipsec-interface-id** statement and the **dedicated** or **shared** statement at the **[edit interfaces interface-name unit logical-unit-number dial-options]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number dial-options]
ipsec-interface-id identifier;
(dedicated | shared);
```

Specifying the interface identifier in the **dial-options** statement makes this logical interface part of the pool identified by the **ipsec-interface-id** statement.



NOTE: Only one interface identifier can be specified at a time. You can include the **ipsec-interface-id** statement or the **l2tp-interface-id** statement, but not both.

If you configure **shared** mode, it enables one logical interface to be shared across multiple tunnels. The **dedicated** statement specifies that the logical interface is used in a dedicated mode, which is necessary when you are configuring an IPsec link-type tunnel. You must include the **dedicated** statement when you specify an **ipsec-interface-id** value.

Default IKE and IPsec Proposals

The software includes implicit default IKE and IPsec proposals to match the proposals sent by the dynamic peers. The values are shown in [Table 24 on page 460](#); if more than one value is shown, the first value is the default.



NOTE: RSA certificates are not supported with dynamic endpoint configuration.

Table 24: Default IKE and IPsec Proposals for Dynamic Negotiations

Statement Name	Values
Implicit IKE Proposal	
authentication-method	pre-shared keys
dh-group	group1, group2, group5, group14
authentication-algorithm	sha1, md5, sha-256
encryption-algorithm	3des-cbc, des-cbc, aes-128, aes-192, aes-256
lifetime-seconds	3600 seconds
Implicit IPsec Proposal	
protocol	esp, ah, bundle
authentication-algorithm	hmac-sha1-96, hmac-md5-96
encryption-algorithm	3des-cbc, des-cbc, aes-128, aes-192, aes-256
lifetime-seconds	28,800 seconds (8 hours)

Related Documentation

- [Configuring IKE Policies on page 409](#)
- [Configuring IPsec Rules on page 422](#)
- [Configuring IKE Proposals on page 405](#)
- [Configuring IPsec Proposals on page 415](#)
- [Configuring Security Associations on page 385](#)

Example: Configuring Dynamically Assigned Policy Based Tunnels

This example shows how to configure dynamically assigned policy-based tunnels and contains the following sections.

- [Requirements on page 461](#)
- [Overview and Topology on page 461](#)
- [Configuration on page 462](#)
- [Verification on page 466](#)

Requirements

This example uses the following hardware and software components:

- Three M Series, MX Series or T Series routers.
- Junos OS Release 9.4 or later.

Overview and Topology

An IPsec policy for dynamic endpoints defines a combination of security parameters (IPsec proposals) used during IPsec negotiation between dynamic peer security gateways, in which the remote ends of tunnels do not have a statically assigned IP address.


A policy based VPN is a configuration with a specific VPN tunnel referenced in a policy which acts as a Tunnel. You use a Policy-based VPN if the remote VPN device is a non-Juniper device and if you must access only one subnet or one network at the remote site, across the VPN.

This example explains the IPsec dynamic endpoint tunneling topology as shown in [Figure 22 on page 462](#).

Before you configure dynamically assigned tunnels, be sure you have:

- A local network N-1 connected to a security gateway SG-1. The exit points must have a Juniper Networks router to terminate the static and dynamic peer endpoints. The tunnel termination address on SG-1 is 10.1.1.1 and the local network address is 172.16.1.0/24.
- Two remote peer routers that obtain addresses from an ISP pool and run an RFC-compliant IKE. The remote network N-2 has the address 172.16.2.0/24 and is connected to the security gateway SG-2 with the tunnel termination address 10.2.2.2. The remote network N-3 has the address 172.16.3.0/24 and is connected to the security gateway SG-3 with the tunnel termination address 10.3.3.3.

Figure 22: IPsec Dynamic Endpoint Tunneling Topology



Configuration

To configure dynamically assigned policy based tunnels, perform these tasks:



NOTE: The interface types shown in this example are for indicative purpose only. For example, you can use so- interfaces instead of ge- and sp- instead of ms-.

- [Configuring a Next-Hop SGI Service-Set on page 463](#)
- [Results on page 464](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI, at the [edit] hierarchy level, of SGI router.

Configuring Interfaces

```
set interfaces ms-0/0/0 unit 0 family inet
set interfaces ms-0/0/0 unit 1 family inet
set interfaces ms-0/0/0 unit 1 service-domain inside
set interfaces ms-0/0/0 unit 1 dial-options ipsec-interface-id demo-ipsec-interface-id
set interfaces ms-0/0/0 unit 1 dial-options mode shared
set interfaces ms-0/0/0 unit 2 family inet
set interfaces ms-0/0/0 unit 2 service-domain outside
```

Configuring Access Profile

```
set access profile demo-access-profile client * ike allowed-proxy-pair remote 172.16.2.0/24
local 172.16.1.0/24
set access profile demo-access-profile client * ike allowed-proxy-pair remote 172.16.3.0/24
local 172.16.1.0/24
set access profile demo-access-profile client * ike ascii-text keyfordynamicpeers
set access profile demo-access-profile client * ike interface-id demo-ipsec-interface-id
```

Configuring Service Set

```
set services service-set demo-service-set next-hop-service inside-service-interface
ms-0/0/0.1
set services service-set demo-service-set next-hop-service outside-service-interface
ms-0/0/0.2
```


Configuring IPsec Properties	<pre> set services ipsec-vpn ipsec proposal ipsec_proposal_demo1 protocol esp set services ipsec-vpn ipsec proposal ipsec_proposal_demo1 authentication-algorithm hmac-sha1-96 set services ipsec-vpn ipsec proposal ipsec_proposal_demo1 encryption-algorithm 3des-cbc set services ipsec-vpn ipsec policy demo2 perfect-forward-secrecy keys group2 set services ipsec-vpn ipsec policy demo2 proposals ipsec_proposal_demo1 set services ipsec-vpn ike proposal ike_proposal_demo1 authentication-method pre-shared-keys set services ipsec-vpn ike proposal ike_proposal_demo1 dh-group group2 set services ipsec-vpn ike policy ike_policy_demo1 version 2 set services ipsec-vpn ike policy ike_policy_demo1 proposals ike_proposal_demo1 set services ipsec-vpn ike policy ike_policy_demo1 pre-shared-key ascii-text keyfordemo1 </pre>
Configuring Routing Instances	<pre> set routing-instances demo-vrf instance-type vrf set routing-instances demo-vrf ms-0/0/0.1 set routing-instances demo-vrf ms-0/0/0.2 </pre>

Configuring a Next-Hop SGI Service-Set

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy.

1. Configure the interfaces.


```

[edit interfaces]
user@router1# set interfaces ms-0/0/0 unit 0 family inet
user@router1# set interfaces ms-0/0/0 unit 1 family inet
user@router1# set interfaces ms-0/0/0 unit 1 service-domain inside
user@router1# set interfaces ms-0/0/0 unit 1 dial-options ipsec-interface-id
  demo-ipsec-interface-id
user@router1# set interfaces ms-0/0/0 unit 1 dial-options mode shared
user@router1# set interfaces ms-0/0/0 unit 2 family inet
user@router1# set interfaces ms-0/0/0 unit 2 service-domain outside

```
2. Configure the access profile.


```

[edit access]
user@router1# set profile demo-access-profile client * ike allowed-proxy-pair remote
  172.16.2.0/24 local 172.16.1.0/24
user@router1# set profile demo-access-profile client * ike ascii-text
  keyfordynamicpeers
user@router1# set profile demo-access-profile client * ike interface-id
  demo-ipsec-interface-id

```
3. Configure the services set.


```

[edit services]
user@router1# set service-set demo-service-set next-hop-service
  inside-service-interface ms-0/0/0.1
user@router1# set service-set demo-service-set next-hop-service
  outside-service-interface ms-0/0/0.2

```
4. Configure the IPsec properties.


```

[edit services ipsec-vpn]
user@router1# set ipsec proposal ipsec_proposal_demo1 protocol esp

```

```
user@router1#set ipsec proposal ipsec_proposal_demo1 authentication-algorithm
hmac-sha1-96
user@router1#set ipsec proposal ipsec_proposal_demo1 encryption-algorithm
3des-cbc
user@router1#set ipsec policy demo2 perfect-forward-secrecy keys group2
user@router1#set ipsec policy demo2 proposals ipsec_proposal_demo1
user@router1#set ike proposal ike_proposal_demo1 authentication-method
pre-shared-keys
user@router1#set ike proposal ike_proposal_demo1 dh-group group2
user@router1#set ike policy ike_policy_demo1 version 2
user@router1#set ike policy ike_policy_demo1 proposals ike_proposal_demo1
user@router1#set ike policy ike_policy_demo1 pre-shared-key ascii-text keyfordemo1
```

5. Configure the routing instances.

```
[edit routing-instances]
user@router1# set demo-vrf instance-type vrf
user@router1# set demo-vrf ms-0/0/0.1
user@router1# set demo-vrf ms-0/0/0.2
```

Results

From configuration mode of Router 1, confirm your configuration by entering the **show interfaces**, **show access**, and **show services** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
interfaces {
  ms-0/0/0 {
    unit 0 {
      family inet;
    }
    unit 1 {
      family inet;
      service-domain inside;
      dial-options {
        ipsec-interface-id demo-ipsec-interface-id;
        mode shared;
      }
    }
    unit 2 {
      family inet;
      service-domain outside;
    }
  }
}
access {
  profile demo-access-profile client * {
    ike {
      allowed-proxy-pair {
        remote 172.16.2.0/24 local 172.16.1.0/24; #Set for Network 2 connected to Network
        1
        remote 172.16.3.0/24 local 172.16.1.0/24; #Set for Network 3 connected to Network
        1
      }
      pre-shared-key {
        ascii-text keyfordynamicpeers;
      }
    }
  }
}
```

```

    }
    interface-id demo-ipsec-interface-id;
  }
}
services {
  service-set demo-service-set {
    next-hop-service {
      inside-service-interface ms-0/0/0.1;
      outside-service-interface ms-0/0/0.2;
    }
    ipsec-vpn-options {
      local-gateway 1.1.1.1;
      ike-access-profile demo-access-profile;
    }
  }
}
ipsec-vpn {
  ipsec {
    proposal ipsec_proposal_demo1 {
      protocol esp;
      authentication-algorithm hmac-sha1-96;
      encryption-algorithm 3des-cbc;
    }
    policy demo2 {
      perfect-forward-secrecy {
        keys group2;
      }
      proposals ipsec_proposal_demo1;
    }
  }
  ike {
    proposal ike_proposal_demo1 {
      authentication-method pre-shared-keys;
      dh-group group2;
    }
    policy ike_policy_demo1 {
      version 2;
      proposals ike_proposal_demo1;
      pre-shared-key ascii-text "$9$jokmT69pRhrz3hrev7Nik"; ## SECRET-DATA
    }
  }
}
}
routing-instances {
  demo-vrf {
    instance-type vrf;
    interface ms-0/0/0.1;
    interface ms-0/0/0.2;
  }
}
}

```

Verification

Verifying That the Next-Hop SGI Service Set with Policy-Based Tunnels Is Created

Purpose Verify that the next-hop SGI service set with policy-based tunnels is created.

Action From operational mode, enter the **show route** command.

```
user@router1> show route
demo-vrf.inet.0: .... # Routing instance
172.11.0.0/24 *[Static/1]..
> via ms-0/0/0.1
172.12.0.0/24 *[Static/1]..
> via ms-0/0/0.1
```

From operational mode, enter the **show services ipsec-vpn ipsec security-associations detail**

```
user@router1>show services ipsec-vpn ipsec security-associations detail
rule: junos-dynamic-rule-0
term: term-0
local-gateway-address : 10.1.1.1 #Tunnel termination address on SG-1
remote-gateway-address: 10.2.2.2 #Tunnel termination address on SG-2
source-address : 0.0.0.0/0
destination-address : 0.0.0.0/0
ipsec-inside-interface: ms-0/0/0.1
term: term-1
local-gateway-address : 10.1.1.1 #Tunnel termination address on SG-1
remote-gateway-address: 10.3.3.3 #Tunnel termination address on SG-3
source-address : 0.0.0.0/0
destination-address : 0.0.0.0/0
IPsec Properties
ipsec-inside-interface: ms-0/0/0.1
match-direction: input
```

Meaning The **show services ipsec-vpn ipsec security-associations detail** command output shows the properties that you configured.

Related Documentation

- [Understanding Junos VPN Site Secure on page 371](#)
- [Configuring Security Associations on page 385](#)
- [Configuring IPsec Policies on page 420](#)
- [Configuring IKE Policies on page 409](#)
- [Tracing Junos VPN Site Secure Operations on page 436](#)

Example: Configuring IKE Dynamic SAs

This example shows how to configure IKE dynamic SAs and contains the following sections.

- [Requirements on page 467](#)
- [Overview and Topology on page 467](#)

- [Configuration on page 468](#)
- [Verification on page 479](#)

Requirements

This example uses the following hardware and software components:

- Four M Series, MX Series, or T Series routers with multiservices interfaces installed in them.
- Junos OS Release 9.4 or later.

No special configuration beyond device initiation is required before you can configure this feature.

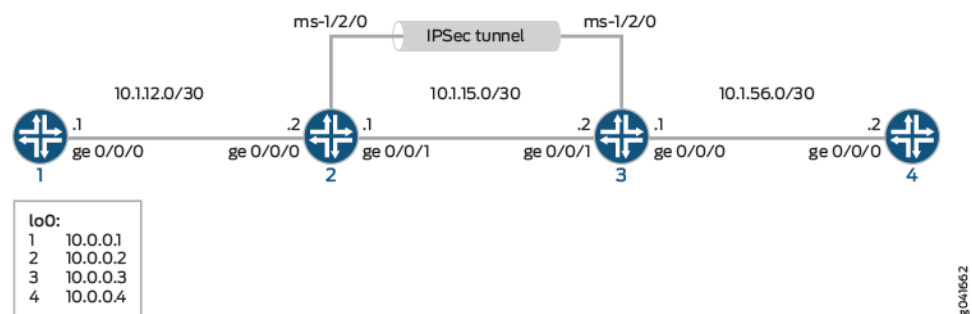
Overview and Topology

A security association (SA) is a simplex connection that enables two hosts to securely communicate with each other by means of IPsec.

Dynamic SAs are best suited for large-scale, geographically distributed networks where manual distribution, maintenance, and tracking of keys are difficult tasks. Dynamic SAs are configured with a set of proposals that are negotiated by the security gateways. The keys are generated as part of the negotiation and do not need to be specified in the configuration. A dynamic SA includes one or more proposals that allow you to prioritize a list of protocols and algorithms to be negotiated with the peer.

[Figure 23 on page 467](#) shows an IPsec topology that contains a group of four routers. This configuration requires Routers 2 and 3 to establish an IPsec tunnel by using an IKE dynamic SA, enhanced authentication, and encryption. Routers 1 and 4 provide basic connectivity and are used to verify that the IPsec tunnel is operational.

Figure 23: IKE Dynamic SAs



NOTE: When you do not specify an IKE proposal, an IPsec proposal, and an IPsec policy on a MultiServices PIC, the Junos OS defaults to the highest level of encryption and authentication. As a result, the default authentication protocol is ESP, the default authentication mode is HMAC-SHA1-96, and the default encryption mode is 3DES-CBC.

Configuration

To configure IKE dynamic SA, perform these tasks:



NOTE: The interface types shown in this example are for indicative purpose only. For example, you can use `so-` interfaces instead of `ge-` and `sp-` instead of `ms-`.

- [Configuring Router 1 on page 468](#)
- [Configuring Router 2 on page 469](#)
- [Configuring Router 3 on page 474](#)
- [Configuring Router 4 on page 478](#)

Configuring Router 1

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI, at the [edit] hierarchy level, of Router 1.

```
set interfaces ge-0/0/0 description "to R2 ge-0/0/0"
set interfaces ge-0/0/0 unit 0 family inet address 10.1.12.2/30
set interfaces lo0 unit 0 family inet address 10.0.0.1/32
set routing-options router-id 10.0.0.1
set protocols ospf area 0.0.0.0 interface ge-0/0/0
set protocols ospf area 0.0.0.0 interface lo0.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router 1 for OSPF connectivity with Router 2:

1. Configure an Ethernet interface and a loopback interface.


```
[edit interfaces]
user@router1# set ge-0/0/0 description "to R2 ge-0/0/0"
user@router1# set ge-0/0/0 unit 0 family inet address 10.1.12.2/30
user@router1# set lo0 unit 0 family inet address 10.0.0.1/32
```
2. Specify the OSPF area and associate the interfaces with the OSPF area.


```
[edit interfaces]
user@router1# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@router1# set ospf area 0.0.0.0 interface lo0.0
```
3. Configure the router ID.


```
[edit routing-options]
user@router1# set router-id 10.0.0.1
```
4. Commit the configuration.

```
[edit]
user@router1# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration

```
user@router1# show interfaces
interfaces {
  ge-0/0/0 {
    description "To R2 ge-0/0/0";
    unit 0 {
      family inet {
        address 10.1.12.2/30;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.1/32;
      }
    }
  }
}

user@router1# show protocols ospf
protocols {
  ospf {
    area 0.0.0.0 {
      interface ge-0/0/0.0;
      interface lo0.0;
    }
  }
}

user@router1# show routing-options
routing-options {
  router-id 10.0.0.1;
}
```

Configuring Router 2

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI, at the [edit] hierarchy level, of Router 2.

```
set interfaces ge-0/0/0 description "to R1 ge-0/0/0"
set interfaces ge-0/0/0 unit 0 family inet address 10.1.12.1/30
set interfaces ge-0/0/1 description "to R3 ge-0/0/1"
set interfaces ge-0/0/1 unit 0 family inet address 10.1.15.1/30
set interfaces ms-1/2/0 services-options syslog host local services info
set interfaces ms-1/2/0 unit 0 family inet
```

```

set interfaces ms-1/2/0 unit 1 family inet
set interfaces ms-1/2/0 unit 1 service-domain inside
set interfaces ms-1/2/0 unit 2 family inet
set interfaces ms-1/2/0 unit 2 service-domain outside
set interfaces lo0 unit 0 family inet address 10.0.0.2/32
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ms-1/2/0.1
set routing-options router-id 10.0.0.2
set services ipsec-vpn rule rule-ike term term-ike then remote-gateway 10.1.15.2
set services ipsec-vpn rule rule-ike term term-ike then dynamic ike-policy ike-demo-policy
set services ipsec-vpn rule rule-ike term term-ike then dynamic ipsec-policy
    ipsec-demo-policy
set services ipsec-vpn rule match-direction input
set services ipsec-vpn ike proposal ike-demo-proposal authentication-method
    pre-shared-keys
set services ipsec-vpn ike proposal ike-demo-proposal dh-group group2
set services ipsec-vpn ike policy ike-demo-policy pre-shared proposals demo-proposal
set services ipsec-vpn ike policy ike-demo-policy pre-shared pre-shared-key ascii-text
    keyfordemo
set services ipsec-vpn ipsec proposal ipsec-demo-proposal protocol esp
set services ipsec-vpn ipsec proposal ipsec-demo-proposal authentication-algorithm
    hmac-sha1-96
set services ipsec-vpn ipsec proposal ipsec-demo-proposal encryption-algorithm 3des-cbc
set services ipsec-vpn ipsec policy ipsec-demo-policy perfect-forward-secrecy keys
    group2
set services ipsec-vpn ipsec proposals ipsec-demo-proposal
set services service-set demo-service-set next-hop-service inside-service-interface
    ms-1/2/0.1
set services service-set demo-service-set next-hop-service outside-service-interface
    ms-1/2/0.2
set services service-set demo-service-set ipsec-vpn-options local-gateway 10.1.15.1
set services service-set demo-service-set ipsec-vpn-rules rule-ike

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure OSPF connectivity and IPsec tunnel parameters on Router 2:

1. Configure interface properties. In this step, you configure two Ethernet interfaces (ge-1/0/0 and ge-1/0/1), a loopback interface, and a multiservices interface (ms-1/2/0).

```

[edit interfaces]
user@router2# set ge-0/0/0 description "to R1 ge-0/0/0"
user@router2# set ge-0/0/0 unit 0 family inet address 10.1.12.1/30
user@router2# set ge-0/0/1 description "to R3 ge-0/0/1"
user@router2# set ge-0/0/1 unit 0 family inet address 10.1.15.1/30
user@router2# set ms-1/2/0 services-options syslog host local services info
user@router2# set ms-1/2/0 unit 0 family inet
user@router2# set ms-1/2/0 unit 1 family inet
user@router2# set ms-1/2/0 unit 1 service-domain inside
user@router2# set ms-1/2/0 unit 2 family inet
user@router2# set ms-1/2/0 unit 2 service-domain outside

```



```
user@router2# set lo0 unit 0 family inet address 10.0.0.2/32
```

2. Specify the OSPF area and associate the interfaces with the OSPF area.

```
[edit protocols]
```

```
user@router2# set ospf area 0.0.0.0 interface ge-0/0/0.0
```

```
user@router2# set ospf area 0.0.0.0 interface lo0.0
```

```
user@router2# set ospf area 0.0.0.0 interface ms-1/2/0.1
```

3. Configure the router ID.

```
[edit routing-options]
```

```
user@router2# set router-ID 10.0.0.2
```

4. Configure an IPsec rule. In this step, you configure an IPsec rule, specify manual SA parameters, such as the remote gateway address, authentication and encryption properties, and so on.



NOTE: By default, Junos OS uses IKE policy version 1.0. Junos OS Release 11.4 and later also support IKE policy version 2.0 which you must configure at [edit services ipsec-vpn ike policy *policy-name* pre-shared].

```
[edit services ipsec-vpn]
```

```
user@router2# set rule rule-ike term term-ike then remote-gateway 10.1.15.2
```

```
user@router2# set rule rule-ike term term-ike then dynamic ike-policy  
ike-demo-policy
```

```
user@router2# set rule rule-ike term term-ike then dynamic ipsec-policy  
ipsec-demo-policy
```

```
user@router2# set rule match-direction input
```

```
user@router2# set ike proposal ike-demo-proposal authentication-method  
pre-shared-keys
```

```
user@router2# set ike proposal ike-demo-proposal dh-group group2
```

```
user@router2# set ike policy ike-demo-policy pre-shared proposals demo-proposal
```

```
user@router2# set ike policy ike-demo-policy pre-shared pre-shared-key ascii-text  
keyfordemo
```

```
user@router2# set ipsec proposal ipsec-demo-proposal protocol esp
```

```
user@router2# set ipsec proposal ipsec-demo-proposal authentication-algorithm  
hmac-sha1-96
```

```
user@router2# set ipsec proposal ipsec-demo-proposal encryption-algorithm  
3des-cbc
```

```
user@router2# set ipsec policy ipsec-demo-policy perfect-forward-secrecy keys  
group2
```

```
user@router2# set ipsec proposals ipsec-demo-proposal
```

5. Configure a next-hop style service set, specify the local-gateway address, and associate the IPsec VPN rule with the service set.

```
[edit services]
```

```
user@router2# set service-set demo-service-set next-hop-service  
inside-service-interface ms-1/2/0.1
```

```
user@router2# set service-set demo-service-set next-hop-service  
outside-service-interface ms-1/2/0.2
```

```
user@router2# set service-set demo-service-set ipsec-vpn-options local-gateway  
10.1.15.1
```

```
user@router2# set service-set demo-service-set ipsec-vpn-rules rule-ike
```

6. Commit the configuration.

```
[edit]
user@router2# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show routing-options**, and **show services** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration

```
user@router1# show interfaces
interfaces {
  ge-0/0/0 {
    description "To R1 ge-0/0/0";
    unit 0 {
      family inet {
        address 10.1.12.1/30;
      }
    }
  }
  ge-0/0/1 {
    description "To R3 ge-0/0/1";
    unit 0 {
      family inet {
        address 10.1.15.1/30;
      }
    }
  }
  ms-1/2/0 {
    services-options {
      syslog {
        host local {
          services info;
        }
      }
    }
    unit 0 {
      family inet;
    }
    unit 1 {
      family inet;
      service-domain inside;
    }
    unit 2 {
      family inet;
      service-domain outside;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.2/32;
      }
    }
  }
}
```

```

}

user@router2# show protocols ospf
protocols {
  ospf {
    area 0.0.0.0 {
      interface ge-0/0/0.0;
      interface lo0.0;
      interface ms-1/2/0.1;
    }
  }
}

user@router2# show routing-options
routing-options {
  router-id 10.0.0.2;
}

user@router2# show services
services {
  ipsec-vpn {
    rule rule-ike {
      term term-ike {
        then {
          remote-gateway 10.1.15.2;
          dynamic {
            ike-policy ike-demo-policy;
            ipsec-policy ipsec-demo-policy;
          }
        }
      }
    }
    match-direction input;
  }
  ike {
    proposal ike-demo-proposal {
      authentication-method pre-shared-keys;
      dh-group group2;
    }
    policy ike-demo-policy {
      proposals demo-proposal;
      pre-shared-key ascii-text "$9$jokmT69pRhrz3hrev7Nik"; ## SECRET-DATA
    }
  }
  ipsec {
    proposal ipsec-demo-proposal {
      protocol esp;
      authentication-algorithm hmac-sha1-96;
      encryption-algorithm 3des-cbc;
    }
    policy ipsec-demo-policy {
      perfect-forward-secrecy {
        keys group2;
      }
      proposals ipsec-demo-proposal;
    }
  }
}

```

```
service-set demo-service-set {
  next-hop-service {
    inside-service-interface ms-1/2/0.1;
    outside-service-interface ms-1/2/0.2;
  }
  ipsec-vpn-options {
    local-gateway 10.1.15.1;
  }
  ipsec-vpn-rules rule-ike;
}
service-set demo-service-set {
  next-hop-service {
    inside-service-interface ms-1/2/0.1;
    outside-service-interface ms-1/2/0.2;
  }
  ipsec-vpn-options {
    local-gateway 10.1.15.2;
  }
  ipsec-vpn-rules rule-ike;
}
```

Configuring Router 3

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI, at the [edit] hierarchy level, of Router 3.

```
set interfaces ge-0/0/0 description "to R4 ge-0/0/0"
set interfaces ge-0/0/0 unit 0 family inet address 10.1.56.1/30
set interfaces ge-0/0/1 description "to R2 ge-0/0/1"
set interfaces ge-0/0/1 unit 0 family inet address 10.1.15.2/30
set interfaces ms-1/2/0 services-options syslog host local services info
set interfaces ms-1/2/0 unit 0 family inet
set interfaces ms-1/2/0 unit 1 family inet
set interfaces ms-1/2/0 unit 1 service-domain inside
set interfaces ms-1/2/0 unit 2 family inet
set interfaces ms-1/2/0 unit 2 service-domain outside
set interfaces lo0 unit 0 family inet address 10.0.0.3/32
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ms-1/2/0.1
set routing-options router-id 10.0.0.3
set services ipsec-vpn rule rule-ike term term-ike then remote-gateway 10.1.15.1
set services ipsec-vpn rule rule-ike term term-ike then dynamic ike-policy ike-demo-policy
set services ipsec-vpn rule rule-ike term term-ike then dynamic ipsec-policy
  ipsec-demo-policy
set services ipsec-vpn rule match-direction input
set services ipsec-vpn ike proposal ike-demo-proposal authentication-method
  pre-shared-keys
set services ipsec-vpn ike proposal ike-demo-proposal dh-group group2
set services ipsec-vpn ike policy ike-demo-policy pre-shared proposals demo-proposal
set services ipsec-vpn ike policy ike-demo-policy pre-shared pre-shared-key ascii-text
  keyfordemo
set services ipsec-vpn ipsec proposal ipsec-demo-proposal protocol esp
```

```

set services ipsec-vpn ipsec proposal ipsec-demo-proposal authentication-algorithm
  hmac-sha1-96
set services ipsec-vpn ipsec proposal ipsec-demo-proposal encryption-algorithm 3des-cbc
set services ipsec-vpn ipsec policy ipsec-demo-policy perfect-forward-secrecy keys
  group2
set services ipsec-vpn ipsec proposals ipsec-demo-proposal
set services service-set demo-service-set next-hop-service inside-service-interface
  ms-1/2/0.1
set services service-set demo-service-set next-hop-service outside-service-interface
  ms-1/2/0.2
set services service-set demo-service-set ipsec-vpn-options local-gateway 10.1.15.2
set services service-set demo-service-set ipsec-vpn-rules rule-ike

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure OSPF connectivity and IPsec tunnel parameters on Router 3:

1. Configure interface properties. In this step, you configure two Ethernet interfaces (ge-1/0/0 and ge-1/0/1), a loopback interface, and a multiservices interface (ms-1/2/0).

```

[edit interfaces]
user@router3# set ge-0/0/0 description "to R4 ge-0/0/0"
user@router3# set ge-0/0/0 unit 0 family inet address 10.1.56.1/30
user@router3# set ge-0/0/1 description "to R2 ge-0/0/1"
user@router3# set ge-0/0/1 unit 0 family inet address 10.1.15.2/30
user@router3# set ms-1/2/0 services-options syslog host local services info
user@router3# set ms-1/2/0 unit 0 family inet
user@router3# set ms-1/2/0 unit 1 family inet
user@router3# set ms-1/2/0 unit 1 service-domain inside
user@router3# set ms-1/2/0 unit 2 family inet
user@router3# set ms-1/2/0 unit 2 service-domain outside
user@router3# set lo0 unit 0 family inet address 10.0.0.3/32

```

2. Specify the OSPF area and associate the interfaces with the OSPF area.

```

[edit protocols]
user@router3# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@router3# set ospf area 0.0.0.0 interface lo0.0
user@router3# set ospf area 0.0.0.0 interface ms-1/2/0.1

```

3. Configure a router ID.

```

[edit routing-options]
user@router3# set router-id 10.0.0.3

```

4. Configure an IPsec rule. In this step, you configure an IPsec rule and specify manual SA parameters, such as the remote gateway address, authentication and encryption properties, and so on.

```

[edit services ipsec-vpn]
user@router3# set rule rule-ike term term-ike then remote-gateway 10.1.15.1
user@router3# set rule rule-ike term term-ike then dynamic ike-policy
  ike-demo-policy

```

```

user@router3# set rule rule-ike term term-ike then dynamic ipsec-policy
ipsec-demo-policy
user@router3# set rule match-direction input
user@router3# set ike proposal ike-demo-proposal authentication-method
pre-shared-keys
user@router3# set ike proposal ike-demo-proposal dh-group group2
user@router3# set ike policy ike-demo-policy pre-shared proposals demo-proposal
user@router3# set ike policy ike-demo-policy pre-shared pre-shared-key ascii-text
keyfordemo
user@router3# set ipsec proposal ipsec-demo-proposal protocol esp
user@router3# set ipsec proposal ipsec-demo-proposal authentication-algorithm
hmac-sha1-96
user@router3# set ipsec proposal ipsec-demo-proposal encryption-algorithm
3des-cbc
user@router3# set ipsec policy ipsec-demo-policy perfect-forward-secrecy keys
group2
user@router3# set ipsec proposals ipsec-demo-proposal

```

5. Configure a next-hop style service set, specify the local-gateway address, and associate the IPsec VPN rule with the service set.

```

[edit services]
user@router3# set service-set demo-service-set next-hop-service
inside-service-interface ms-1/2/0.1
user@router3# set service-set demo-service-set next-hop-service
outside-service-interface ms-1/2/0.2
user@router3# set service-set demo-service-set ipsec-vpn-options local-gateway
10.1.15.2
user@router3# set service-set demo-service-set ipsec-vpn-rules rule-ike

```

6. Commit the configuration.

```

[edit]
user@router3# commit

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show routing-options**, and **show services** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration

```

user@router3# show interfaces
interfaces {
  ge-0/0/0 {
    description "To R4 ge-0/0/0";
    unit 0 {
      family inet {
        address 10.1.56.1/30;
      }
    }
  }
  ge-0/0/1 {
    description "To R2 ge-0/0/1";
    unit 0 {
      family inet {
        address 10.1.15.2/30;
      }
    }
  }
}

```

```

    }
  }
  ms-1/2/0 {
    services-options {
      syslog {
        host local {
          services info;
        }
      }
    }
  }
  unit 0 {
    family inet {
    }
  }
  unit 1 {
    family inet;
    service-domain inside;
  }
  unit 2 {
    family inet;
    service-domain outside;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.0.0.3/32;
    }
  }
}
}
}

user@router3# show protocols ospf
protocols {
  ospf {
    area 0.0.0.0 {
      interface ge-0/0/0.0;
      interface lo0.0;
      interface ms-1/2/0.1;
    }
  }
}

user@router3# show routing-options
routing-options {
  router-id 10.0.0.3;
}

user@router3# show services
services {
  ipsec-vpn {
    rule rule-ike {
      term term-ike {
        then {
          remote-gateway 10.1.15.1;
          dynamic {
            ike-policy ike-demo-policy;
          }
        }
      }
    }
  }
}

```

```

        ipsec-policy ipsec-demo-policy;
    }
}
match-direction input;
}
ike {
    proposal ike-demo-proposal {
        authentication-method pre-shared-keys;
        dh-group group2;
    }
    policy ike-demo-policy {
        proposals demo-proposal;
        pre-shared-key ascii-text "$9$jokmT69pRhrz3hrev7Nik"; ## SECRET-DATA
    }
}
ipsec {
    proposal ipsec-demo-proposal {
        protocol esp;
        authentication-algorithm hmac-sha1-96;
        encryption-algorithm 3des-cbc;
    }
    policy ipsec-demo-policy {
        perfect-forward-secrecy {
            keys group2;
        }
        proposals ipsec-demo-proposal;
    }
}
}

```

Configuring Router 4

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI, at the [edit] hierarchy level, of Router 4.

```

set interfaces ge-0/0/0 description "to R3 ge-0/0/0"
set interfaces ge-0/0/0 unit 0 family inet address 10.1.56.2/30
set interfaces lo0 unit 0 family inet address 10.0.0.4/32
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set routing-options router-id 10.0.0.4

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To set up OSPF connectivity with Router 4

1. Configure the interfaces. In this step, you configure an Ethernet interface (ge-1/0/1) and a loopback interface.

```

user@router4# set interfaces ge-0/0/0 description "to R3 ge-0/0/0"

```



```
user@router4# set interfaces ge-0/0/0 unit 0 family inet address 10.1.56.2/30
user@router4# set interfaces lo0 unit 0 family inet address 10.0.0.4/32
```

- Specify the OSPF area and associate the interfaces with the OSPF area.

```
user@router4# set protocols ospf area 0.0.0.0 interface ge-0/0/0
user@router4# set protocols ospf area 0.0.0.0 interface lo0.0
```

- Configure the router ID.

```
[edit routing-options]
user@router4# set router-id 10.0.0.4
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration

```
user@router4# show interfaces
interfaces {
  ge-0/0/0 {
    description "To R3 ge-0/0/0";
    unit 0 {
      family inet {
        address 10.1.56.2/30;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.4/32;
      }
    }
  }
}

user@router4# show protocols ospf
protocols {
  ospf {
    area 0.0.0.0 {
      interface ge-0/0/0.0;
      interface lo0.0;
    }
  }
}

user@router4# show routing-options
routing-options {
  router-id 10.0.0.4;
}
```

Verification

Verifying Your Work on Router 1

Purpose Verify proper operation of Router 1.

Action From operational mode, enter **ping 10.1.56.2** command to the ge-0/0/0 interface on Router 4 to send traffic across the IPsec tunnel

```
user@router1>ping 10.1.56.2
PING 10.1.56.2 (10.1.56.2): 56 data bytes
64 bytes from 10.1.56.2: icmp_seq=0 ttl=254 time=1.351 ms
64 bytes from 10.1.56.2: icmp_seq=1 ttl=254 time=1.187 ms
64 bytes from 10.1.56.2: icmp_seq=2 ttl=254 time=1.172 ms
64 bytes from 10.1.56.2: icmp_seq=3 ttl=254 time=1.154 ms
64 bytes from 10.1.56.2: icmp_seq=4 ttl=254 time=1.156 ms
^C
--- 10.1.56.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.154/1.204/1.351/0.074 ms
```

Meaning The output shows that Router 1 is able to reach Router 4 over the IPsec tunnel.

Verifying Your Work on Router 2

Purpose Verify that the IKE SA negotiation is successful.

Action From operational mode, enter the **show services ipsec-vpn ike security-associations** command.

```
user@router2>show services ipsec-vpn ike security-associations
Remote Address State Initiator cookie Responder cookie Exchange type
10.1.15.2 Matured 03075bd3a0000003 4bff26a5c7000003 Main
```

To verify that the IPsec security association is active, issue the **show services ipsec-vpn ipsec security-associations detail** command. Notice that the SA contains the default settings inherent in the MultiServices PIC, such as ESP for the protocol and HMAC-SHA1-96 for the authentication algorithm.

From operational mode, enter the **show services ipsec-vpn ipsec security-associations detail** command.

```
user@router2> show services ipsec-vpn ipsec security-associations detail
Service set: demo-service-set
Rule: rule-ike, Term: term-ike, Tunnel index: 1
Local gateway: 10.1.15.1, Remote gateway: 10.1.15.2
Local identity: ipv4_subnet(any:0,[0..7]=10.1.12.0/24)
Remote identity: ipv4_subnet(any:0,[0..7]=10.1.56.0/24)
Direction: inbound, SPI: 2666326758, AUX-SPI: 0
Mode: tunnel, Type: dynamic, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Soft lifetime: Expires in 26863 seconds
Hard lifetime: Expires in 26998 seconds
Anti-replay service: Enabled, Replay window size: 64
Direction: outbound, SPI: 684772754, AUX-SPI: 0
Mode: tunnel, Type: dynamic, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Soft lifetime: Expires in 26863 seconds
Hard lifetime: Expires in 26998 seconds
Anti-replay service: Enabled, Replay window size: 64
```

To verify that traffic is traveling through the bidirectional IPsec tunnel, issue the **show services ipsec-vpn statistics** command:

From operational mode, enter the **show services ipsec-vpn statistics** command.

```
user@router2> show services ipsec-vpn ipsec statistics
PIC: ms-1/2/0, Service set: demo-service-set
ESP Statistics:
Encrypted bytes: 2248
Decrypted bytes: 2120
Encrypted packets: 27
Decrypted packets: 25
AH Statistics:
Input bytes: 0
Output bytes: 0
Input packets: 0
Output packets: 0
Errors:
AH authentication failures: 0, Replay errors: 0
ESP authentication failures: 0, ESP decryption failures: 0
Bad headers: 0, Bad trailers: 0
```

Meaning The **show services ipsec-vpn ipsec security-associations detail** command output shows the SA properties that you configured.

The **show services ipsec-vpn ipsec statistics** command output shows the traffic flow over the IPsec tunnel.

Verifying Your Work on Router 3

Purpose Verify that the IKE SA negotiation is successful on Router 3.

Action From operational mode, enter the **show services ipsec-vpn ike security-associations** command. To be successful, the SA on Router 3 must contain the same settings you specified on Router 2.

```
user@router3>show services ipsec-vpn ike security-associations
Remote Address State Initiator cookie Responder cookie Exchange type
10.1.15.1 Matured 03075bd3a0000003 4bff26a5c7000003 Main
```

To verify that the IPsec SA is active, issue the **show services ipsec-vpn ipsec security-associations detail** command. To be successful, the SA on Router 3 must contain the same settings you specified on Router 2.

From operational mode, enter the **show services ipsec-vpn ipsec security-associations detail** command.

```
user@router3>show services ipsec-vpn ipsec security-associations detail
Service set: demo-service-set
Rule: rule-ike, Term: term-ike, Tunnel index: 1
Local gateway: 10.1.15.2, Remote gateway: 10.1.15.1
Local identity: ipv4_subnet(any:0,[0..7]=10.1.56.0/24)
Remote identity: ipv4_subnet(any:0,[0..7]=10.1.12.0/24)
Direction: inbound, SPI: 684772754, AUX-SPI: 0
Mode: tunnel, Type: dynamic, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Soft lifetime: Expires in 26598 seconds
```

```
Hard lifetime: Expires in 26688 seconds
Anti-replay service: Enabled, Replay window size: 64
Direction: outbound, SPI: 2666326758, AUX-SPI: 0
Mode: tunnel, Type: dynamic, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Soft lifetime: Expires in 26598 seconds
Hard lifetime: Expires in 26688 seconds
Anti-replay service: Enabled, Replay window size: 64
```

To verify that traffic is traveling through the bidirectional IPsec tunnel, issue the **show services ipsec-vpn statistics** command:

From operational mode, enter the **show services ipsec-vpn ike security-associations** command.

```
user@router3>show services ipsec-vpn ipsec statistics
PIC: ms-1/2/0, Service set: demo-service-set
ESP Statistics:
Encrypted bytes: 2120
Decrypted bytes: 2248
Encrypted packets: 25
Decrypted packets: 27
AH Statistics:
Input bytes: 0
Output bytes: 0
Input packets: 0
Output packets: 0
Errors:
AH authentication failures: 0, Replay errors: 0
ESP authentication failures: 0, ESP decryption failures: 0
Bad headers: 0, Bad trailers: 0
```

Meaning The **show services ipsec-vpn ipsec security-associations detail** command output shows the SA properties that you configured.

The **show services ipsec-vpn ipsec statistics** command output shows the traffic flow over the IPsec tunnel.

Verifying Your Work on Router 4

Purpose Verify that the IKE SA negotiation is successful.

Action From operational mode, enter **ping 10.1.12.2** command to the ge-0/0/0 interface on Router 1 to send traffic across the IPsec tunnel.

```
user@router4>ping 10.1.12.2
PING 10.1.12.2 (10.1.12.2): 56 data bytes
64 bytes from 10.1.12.2: icmp_seq=0 ttl=254 time=1.350 ms
64 bytes from 10.1.12.2: icmp_seq=1 ttl=254 time=1.161 ms
64 bytes from 10.1.12.2: icmp_seq=2 ttl=254 time=1.124 ms
64 bytes from 10.1.12.2: icmp_seq=3 ttl=254 time=1.142 ms
64 bytes from 10.1.12.2: icmp_seq=4 ttl=254 time=1.139 ms
64 bytes from 10.1.12.2: icmp_seq=5 ttl=254 time=1.116 ms
^C
--- 10.1.12.2 ping statistics ---
6 packets transmitted, 6 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.116/1.172/1.350/0.081 ms
```

To confirm that traffic travels through the IPsec tunnel, issue the **traceroute** command to the ge-0/0/0 interface on Router 1. Notice that the physical interface between Routers 2 and 3 is not referenced in the path; traffic enters the IPsec tunnel through the adaptive services IPsec inside interface on Router 3, passes through the loopback interface on Router 2, and ends at the ge-0/0/0 interface on Router 1.

From operational mode, enter the **traceroute 10.1.12.2**.

```
user@router4>traceroute 10.1.12.2
traceroute to 10.1.12.2 (10.1.12.2), 30 hops max, 40 byte packets
 1 10.1.15.2 (10.1.15.2) 0.987 ms 0.630 ms 0.563 ms
 2 10.0.0.2 (10.0.0.2) 1.194 ms 1.058 ms 1.033 ms
 3 10.1.12.2 (10.1.12.2) 1.073 ms 0.949 ms 0.932 ms
```

Meaning The **ping 10.1.12.2** output shows that Router 4 is able to reach Router 1 over the IPsec tunnel.

The **traceroute 10.1.12.2** output shows that traffic travels the IPsec tunnel.

- Related Documentation**
- [Understanding Junos VPN Site Secure on page 371](#)
 - [Configuring Security Associations on page 385](#)
 - [Configuring IKE Proposals on page 405](#)
 - [Configuring IKE Policies on page 409](#)
 - [Example: Configuring Manual SAs on page 391](#)

Example: IKE Dynamic SA Configuration with Digital Certificates

This example shows how to configure IKE dynamic SA with digital certificates and contains the following sections.

- [Requirements on page 483](#)
- [Overview on page 484](#)
- [Configuration on page 484](#)
- [Verification on page 497](#)

Requirements

This example uses the following hardware and software components:

- Four M Series, MX Series, or T Series routers with multiservices interfaces installed in them.
- Junos OS Release 9.4 or later.

Before you configure this example you must request a CA certificate, create a local certificate, and load these digital certificates into the router. For details, see *Requesting for and Installing a Digital Certificates on Your Router*


Overview

A security association (SA) is a simplex connection that enables two hosts to securely communicate with each other using IPsec. This example explains IKE dynamic SA configuration with digital certificates. The use of digital certificates provides additional security to your IKE tunnel. Using default values in the Services PIC, you do not need to configure an IPsec proposal or IPsec policy. However, you must configure an IKE proposal that specifies the use of digital certificates, reference the IKE proposal and local certificate in an IKE policy, and apply the CA profile to the service set.

Figure 24 on page 484 shows an IPsec topology containing a group of four routers. This configuration requires Routers 2 and 3 to establish an IKE-based IPsec tunnel by using digital certificates in place of preshared keys. Routers 1 and 4 provide basic connectivity and are used to verify that the IPsec tunnel is operational.

Topology

Figure 24: MS PIC IKE Dynamic SA Topology Diagram



Configuration

To configure IKE dynamic SA with digital certificates, perform these tasks:



NOTE: The interface types shown in this example are for indicative purpose only. For example, you can use so- interfaces instead of ge- and sp- instead of ms-.

- [Configuring Router 1 on page 484](#)
- [Configuring Router 2 on page 486](#)
- [Configuring Router 3 on page 491](#)
- [Configuring Router 4 on page 496](#)

Configuring Router 1

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI, at the [edit] hierarchy level, of Router 1.

```

set interfaces ge-0/0/0 description "to R2 ge-0/0/0"
set interfaces ge-0/0/0 unit 0 family inet address 10.1.12.2/30
set interfaces lo0 unit 0 family inet address 10.0.0.1/32
set routing-options router-id 10.0.0.1
set protocols ospf area 0.0.0.0 interface ge-0/0/0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router 1 for OSPF connectivity with Router 2:

1. Configure an Ethernet interface and the loopback interface.

```

[edit interfaces]
user@router1# set ge-0/0/0 description "to R2 ge-0/0/0"
user@router1# set ge-0/0/0 unit 0 family inet address 10.1.12.2/30
user@router1# set lo0 unit 0 family inet address 10.0.0.1/32

```
2. Specify the OSPF area and associate the interfaces with the OSPF area.

```

[edit protocols]
user@router1# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@router1# set ospf area 0.0.0.0 interface lo0.0

```
3. Configure the router ID.

```

[edit routing-options]
user@router1# set router-id 10.0.0.1

```
4. Commit the configuration.

```

[edit]
user@router1# commit

```

Results From the configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@router1# show interfaces
interfaces {
  ge-0/0/0 {
    description "To R2 ge-0/0/0";
    unit 0 {
      family inet {
        address 10.1.12.2/30;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.1/32;
      }
    }
  }
}

```

```
    }  
  }  
  
user@router1# show protocols ospf  
protocols {  
  ospf {  
    area 0.0.0.0 {  
      interface ge-0/0/0.0;  
      interface lo0.0;  
    }  
  }  
}  
  
user@router1# show routing-options  
routing-options {  
  router-id 10.0.0.1;  
}
```

Configuring Router 2

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI, at the [edit] hierarchy level, of Router 2.

```
set interfaces ge-0/0/0 description "to R1 ge-0/0/0"  
set interfaces ge-0/0/0 unit 0 family inet address 10.1.12.1/30  
set interfaces ge-0/0/1 description "to R3 ge-0/0/1"  
set interfaces ge-0/0/1 unit 0 family inet address 10.1.15.1/30  
set interfaces ms-1/2/0 services-options syslog host local services info  
set interfaces ms-1/2/0 unit 0 family inet  
set interfaces ms-1/2/0 unit 1 family inet  
set interfaces ms-1/2/0 unit 1 service-domain inside  
set interfaces ms-1/2/0 unit 2 family inet  
set interfaces ms-1/2/0 unit 2 service-domain outside  
set interfaces lo0 unit 0 family inet address 10.0.0.2/32  
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0  
set protocols ospf area 0.0.0.0 interface lo0.0  
set protocols ospf area 0.0.0.0 interface ms-1/2/0.1  
set routing-options router-id 10.0.0.2  
set services ipsec-vpn rule rule-ike term term-ike then remote-gateway 10.1.15.2  
set services ipsec-vpn rule rule-ike term term-ike then dynamic ike-policy  
    ike-digital-certificates  
set services ipsec-vpn rule rule-ike term term-ike then dynamic ipsec-policy  
    ipsec-demo-policy  
set services ipsec-vpn rule match-direction input  
set services ipsec-vpn ike proposal ike-demo-proposal authentication-method  
    rsa-signatures  
set services ipsec-vpn ike policy ike-digital-certificates proposals ike-demo-proposal  
set services ipsec-vpn ike policy ike-digital-certificates local-id fqdn router2.juniper.net  
set services ipsec-vpn ike policy ike-digital-certificates local-certificate local-entrust2  
set services ipsec-vpn ike policy ike-digital-certificates remote-id fqdn router3.juniper.net  
set services ipsec-vpn ipsec proposal ipsec-demo-proposal protocol esp  
set services ipsec-vpn ipsec proposal ipsec-demo-proposal authentication-algorithm  
    hmac-sha1-96  
set services ipsec-vpn ipsec proposal ipsec-demo-proposal encryption-algorithm 3des-cbc
```



```

set services ipsec-vpn ipsec policy ipsec-demo-policy perfect-forward-secrecy keys
group2
set services ipsec-vpn ipsec proposals ipsec-demo-proposal
set services ipsec-vpn establish-tunnels immediately
set services service-set demo-service-set next-hop-service inside-service-interface
ms-1/2/0.1
set services service-set demo-service-set next-hop-service outside-service-interface
ms-1/2/0.2
set services service-set demo-service-set ipsec-vpn-options trusted-ca entrust
set services service-set demo-service-set ipsec-vpn-options local-gateway 10.1.15.1
set services service-set demo-service-set ipsec-vpn-rules rule-ike

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure OSPF connectivity and IPsec tunnel parameters on Router 2:

1. Configure interface properties. In this step, you configure two Ethernet interfaces (ge-1/0/0 and ge-1/0/1), the loopback interface and a multiservices interface (ms-1/2/0).

```

[edit interfaces]
user@router2# set ge-0/0/0 description "to R1 ge-0/0/0"
user@router2# set ge-0/0/0 unit 0 family inet address 10.1.12.1/30
user@router2# set ge-0/0/1 description "to R3 ge-0/0/1"
user@router2# set ge-0/0/1 unit 0 family inet address 10.1.15.1/30
user@router2# set ms-1/2/0 services-options syslog host local services info
user@router2# set ms-1/2/0 unit 0 family inet
user@router2# set ms-1/2/0 unit 1 family inet
user@router2# set ms-1/2/0 unit 1 service-domain inside
user@router2# set ms-1/2/0 unit 2 family inet
user@router2# set ms-1/2/0 unit 2 service-domain outside
user@router2# set lo0 unit 0 family inet address 10.0.0.2/32

```

2. Specify the OSPF area and associate the interfaces with the OSPF area.

```

[edit protocols]
user@router2# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@router2# set ospf area 0.0.0.0 interface lo0.0
user@router2# set ospf area 0.0.0.0 interface ms-1/2/0.1

```

3. Configure the router ID.

```

[edit routing-options]
user@router2# set router-ID 10.0.0.2

```

4. Configure an IKE proposal and policy. To enable an IKE proposal for digital certificates, include the **rsa-signatures** statement at the **[edit services ipsec-vpn ike proposal proposal-name authentication-method]** hierarchy level. To reference the local certificate in the IKE policy, include the **local-certificate** statement at the **[edit services ipsec-vpn ike policy policy-name]** hierarchy level. To identify the CA or RA in the service set, include the **trusted-ca** statement at the **[edit services service-set service-set-name ipsec-vpn-options]** hierarchy level.



NOTE: For information about creating and installing digital certificates, see *Requesting for and Installing a Digital Certificates on Your Router*

```
[edit services ipsec-vpn]
user@router2# set ike proposal ike-demo-proposal authentication-method
rsa-signatures
user@router2# set ike policy ike-digital-certificates proposals ike-demo-proposal
user@router2# set ike policy ike-digital-certificates local-id fqdn router2.juniper.net
user@router2# set ike policy ike-digital-certificates local-certificate local-entrust2
user@router2# set ike policy ike-digital-certificates remote-id fqdn router3.juniper.net
```

5. Configure an IPsec proposal and policy. Also, set the **established-tunnels** knob to **immediately**.

```
[edit services ipsec-vpn]
user@router2# set ipsec proposal ipsec-demo-proposal protocol esp
user@router2# set ipsec proposal ipsec-demo-proposal authentication-algorithm
hmac-sha1-96
user@router2# set ipsec proposal ipsec-demo-proposal encryption-algorithm
3des-cbc
user@router2# set ipsec policy ipsec-demo-policy perfect-forward-secrecy keys
group2
user@router2# set ipsec proposals ipsec-demo-proposal
user@router2# set establish-tunnels immediately
```

6. Configure an IPsec rule.

```
[edit services ipsec-vpn]
user@router2# set rule rule-ike term term-ike then remote-gateway 10.115.2
user@router2# set rule rule-ike term term-ike then dynamic ike-policy
ike-digital-certificates
user@router2# set rule rule-ike term term-ike then dynamic ipsec-policy
ipsec-demo-policy
user@router2# set rule match-direction input
```

7. Configure a next-hop style service set, specify the local-gateway address, and associate the IPsec VPN rule with the service set.

```
[edit services]
user@router2# set service-set demo-service-set next-hop-service
inside-service-interface ms-1/2/0.1
user@router2# set service-set demo-service-set next-hop-service
outside-service-interface ms-1/2/0.2
user@router2# set service-set demo-service-set ipsec-vpn-options trusted-ca
entrust
user@router2# set service-set demo-service-set ipsec-vpn-options local-gateway
10.115.1
user@router2# set service-set demo-service-set ipsec-vpn-rules rule-ike
```

8. Commit the configuration.

```
[edit]
user@router2# commit
```

Results From the configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show routing-options**, and **show services** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration

```

user@router2# show interfaces
interfaces {
  ge-0/0/0 {
    description "To R1 ge-0/0/0";
    unit 0 {
      family inet {
        address 10.1.12.1/30;
      }
    }
  }
  ge-0/0/1 {
    description "To R3 ge-0/0/1";
    unit 0 {
      family inet {
        address 10.1.15.1/30;
      }
    }
  }
  ms-1/2/0 {
    services-options {
      syslog {
        host local {
          services info;
        }
      }
    }
    unit 0 {
      family inet;
    }
    unit 1 {
      family inet;
      service-domain inside;
    }
    unit 2 {
      family inet;
      service-domain outside;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.2/32;
      }
    }
  }
}

user@router2# show protocols ospf
protocols {
  ospf {
    area 0.0.0.0 {

```

```
        interface ge-0/0/0.0;
        interface lo0.0;
        interface ms-1/2/0.1;
    }
}

user@router2# show routing-options
routing-options {
    router-id 10.0.0.2;
}

user@router2# show services
services {
    ipsec-vpn {
        rule rule-ike {
            term term-ike {
                then {
                    remote-gateway 10.1.15.2;
                    dynamic {
                        ike-policy ike-digital-certificates;
                        ipsec-policy ipsec-demo-policy
                    }
                }
            }
        }
        match-direction input;
    }
    ike {
        proposal ike-demo-proposal {
            authentication-method rsa-signatures;
        }
        policy ike-digital-certificates {
            proposals ike-demo-proposal;
            local-id fqdn router2.juniper.net;
            local-certificate local-entrust2;
            remote-id fqdn router3.juniper.net;
        }
    }
    ipsec {
        proposal ipsec-demo-proposal {
            protocol esp;
            authentication-algorithm hmac-sha1-96;
            encryption-algorithm 3des-cbc;
        }
        policy demo-policy {
            perfect-forward-secrecy {
                keys group2;
            }
            proposals ipsec-demo-proposal;
        }
        establish-tunnels immediately;
    }
    service-set service-set-dynamic-demo-service-set {
        next-hop-service {
            inside-service-interface ms-1/2/0.1;
            outside-service-interface ms-1/2/0.2;
        }
    }
}
```

```

        ipsec-vpn-options {
            trusted-ca entrust;
            local-gateway 10.1.15.1;
        }
        ipsec-vpn-rules rule-ike;
    }
}

```

Configuring Router 3

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI, at the [edit] hierarchy level, of Router 3.

```

set interfaces ge-0/0/0 description "to R4 ge-0/0/0"
set interfaces ge-0/0/0 unit 0 family inet address 10.1.56.1/30
set interfaces ge-0/0/1 description "to R2 ge-0/0/1"
set interfaces ge-0/0/1 unit 0 family inet address 10.1.15.2/30
set interfaces ms-1/2/0 services-options syslog host local services info
set interfaces ms-1/2/0 unit 0 family inet
set interfaces ms-1/2/0 unit 1 family inet
set interfaces ms-1/2/0 unit 1 service-domain inside
set interfaces ms-1/2/0 unit 2 family inet
set interfaces ms-1/2/0 unit 2 service-domain outside
set interfaces lo0 unit 0 family inet address 10.0.0.3/32
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ms-1/2/0.1
set routing-options router-id 10.0.0.3
set services ipsec-vpn rule rule-ike term term-ike then remote-gateway 10.1.15.1
set services ipsec-vpn rule rule-ike term term-ike then dynamic ike-policy
    ike-digital-certificates
set services ipsec-vpn rule rule-ike term term-ike then dynamic ipsec-policy
    ipsec-demo-policy
set services ipsec-vpn rule match-direction input
set services ipsec-vpn ike proposal ike-demo-proposal authentication-method
    rsa-signatures
set services ipsec-vpn ike policy ike-digital-certificates proposals ike-demo-proposal
set services ipsec-vpn ike policy ike-digital-certificates local-id fqdn router3.juniper.net
set services ipsec-vpn ike policy ike-digital-certificates local-certificate local-entrust3
set services ipsec-vpn ike policy ike-digital-certificates remote-id fqdn router2.juniper.net
set services ipsec-vpn ipsec proposal ipsec-demo-proposal protocol esp
set services ipsec-vpn ipsec proposal ipsec-demo-proposal authentication-algorithm
    hmac-sha1-96
set services ipsec-vpn ipsec proposal ipsec-demo-proposal encryption-algorithm 3des-cbc
set services ipsec-vpn ipsec policy ipsec-demo-policy perfect-forward-secrecy keys
    group2
set services ipsec-vpn ipsec proposals ipsec-demo-proposal
set services ipsec-vpn establish-tunnels immediately
set services service-set demo-service-set next-hop-service inside-service-interface
    ms-1/2/0.1
set services service-set demo-service-set next-hop-service outside-service-interface
    ms-1/2/0.2

```

```

set services service-set demo-service-set ipsec-vpn-options trusted-ca entrust
set services service-set demo-service-set ipsec-vpn-options local-gateway 10.1.15.2
set services service-set demo-service-set ipsec-vpn-rules rule-ike

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.



NOTE: If the IPsec peers do not have a symmetrical configuration containing all the necessary components, they cannot establish a peering relationship. You need to request a CA certificate, create a local certificate, load these digital certificates into the router, and reference them in your IPsec configuration. For information about digital certification, see *Requesting for and Installing a Digital Certificates on Your Router*

To configure OSPF connectivity and IPsec tunnel parameters on Router 3:

1. Configure interface properties. In this step, you configure two Ethernet interfaces (ge-1/0/0 and ge-1/0/1), the loopback interface, and a multiservices interface (ms-1/2/0).

```

[edit interfaces]
user@router3# set ge-0/0/0 description "to R4 ge-0/0/0"
user@router3# set ge-0/0/0 unit 0 family inet address 10.1.56.1/30
user@router3# set ge-0/0/1 description "to R2 ge-0/0/1"
user@router3# set ge-0/0/1 unit 0 family inet address 10.1.15.2/30
user@router3# set ms-1/2/0 services-options syslog host local services info
user@router3# set ms-1/2/0 unit 0 family inet
user@router3# set ms-1/2/0 unit 1 family inet
user@router3# set ms-1/2/0 unit 1 service-domain inside
user@router3# set ms-1/2/0 unit 2 family inet
user@router3# set ms-1/2/0 unit 2 service-domain outside
user@router3# set lo0 unit 0 family inet address 10.0.0.3/32

```

2. Specify the OSPF area, associate the interfaces with the OSPF area.

```

[edit protocols]
user@router3# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@router3# set ospf area 0.0.0.0 interface lo0.0
user@router3# set ospf area 0.0.0.0 interface ms-1/2/0.1

```

3. Configure a router ID.

```

[edit routing-options]
user@router3# set router-id 10.0.0.3

```

4. Configure an IKE proposal and policy. To enable an IKE proposal for digital certificates, include the **rsa-signatures** statement at the **[edit services ipsec-vpn ike proposal proposal-name authentication-method]** hierarchy level. To reference the local certificate in the IKE policy, include the **local-certificate** statement at the **[edit services ipsec-vpn ike policy policy-name]** hierarchy level. To identify the CA or RA in the service set, include the **trusted-ca** statement at the **[edit services service-set service-set-name ipsec-vpn-options]** hierarchy level.



NOTE: For information about creating and installing digital certificates, see *Requesting for and Installing a Digital Certificates on Your Router*

```
[edit services ipsec-vpn]
user@router3# set ike proposal ike-demo-proposal authentication-method
rsa-signatures
user@router3# set ike policy ike-digital-certificates proposals ike-demo-proposal
user@router3# set ike policy ike-digital-certificates local-id fqdn router2.juniper.net
user@router3# set ike policy ike-digital-certificates local-certificate local-entrust2
user@router3# set ike policy ike-digital-certificates remote-id fqdn router3.juniper.net
```

5. Configure an IPsec proposal. Also, set the **established-tunnels** knob to **immediately**.

```
[edit services ipsec-vpn]
user@router3# set ipsec proposal ipsec-demo-proposal protocol esp
user@router3# set ipsec proposal ipsec-demo-proposal authentication-algorithm
hmac-sha1-96
user@router3# set ipsec proposal ipsec-demo-proposal encryption-algorithm
3des-cbc
user@router3# set ipsec policy ipsec-demo-policy perfect-forward-secrecy keys
group2
user@router3# set ipsec proposals ipsec-demo-proposal
user@router3# set establish-tunnels immediately
```

6. Configure an IPsec rule.

```
[edit services ipsec-vpn]
user@router3# set rule rule-ike term term-ike then remote-gateway 10.1.15.2
user@router3# set rule rule-ike term term-ike then dynamic ike-policy
ike-digital-certificates
user@router3# set rule rule-ike term term-ike then dynamic ipsec-policy
ipsec-demo-policy
user@router3# set rule match-direction input
```

7. Configure a next-hop style service set, specify the local-gateway address, and associate the IPsec VPN rule with the service set.

```
[edit services]
user@router3# set service-set demo-service-set next-hop-service
inside-service-interface ms-1/2/0.1
user@router3# set service-set demo-service-set next-hop-service
outside-service-interface ms-1/2/0.2
user@router3# set service-set demo-service-set ipsec-vpn-options trusted-ca
entrust
user@router3# set service-set demo-service-set ipsec-vpn-options local-gateway
10.1.15.2
user@router3# set service-set demo-service-set ipsec-vpn-rules rule-ike
```

8. Commit the configuration.

```
[edit]
user@router3# commit
```

Results From the configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show routing-options**, and **show services** commands. If the output

does not display the intended configuration, repeat the instructions in this example to correct the configuration

```
user@router3# show interfaces
interfaces {
  ge-0/0/0 {
    description "To R4 ge-0/0/0";
    unit 0 {
      family inet {
        address 10.1.56.1/30;
      }
    }
  }
  ge-0/0/1 {
    description "To R2 ge-0/0/1";
    unit 0 {
      family inet {
        address 10.1.15.2/30;
      }
    }
  }
  ms-1/2/0 {
    services-options {
      syslog {
        host local {
          services info;
        }
      }
    }
    unit 0 {
      family inet {
      }
    }
    unit 1 {
      family inet;
      service-domain inside;
    }
    unit 2 {
      family inet;
      service-domain outside;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.3/32;
      }
    }
  }
}

user@router3# show protocols ospf
protocols {
  ospf {
    area 0.0.0.0 {
      interface ge-0/0/0.0;
    }
  }
}
```



```

        interface lo0.0;
        interface ms-1/2/0.1;
    }
}

user@router3# show routing-options
routing-options {
    router-id 10.0.0.3;
}

user@router3# show services
services {
    ipsec-vpn {
        rule rule-ike {
            term term-ike {
                then {
                    remote-gateway 10.1.15.1;
                    dynamic {
                        ike-policy ike-digital-certificates;
                        ipsec-policy ipsec-demo-policy
                    }
                }
            }
        }
        match-direction input;
    }
    ike {
        proposal ike-demo-proposal {
            authentication-method rsa-signatures;
        }
        policy ike-digital-certificates {
            proposals ike-demo-proposal;
            local-id fqdn router3.juniper.net;
            local-certificate local-entrust3;
            remote-id fqdn router2.juniper.net;
        }
    }
    ipsec {
        proposal ipsec-demo-proposal {
            protocol esp;
            authentication-algorithm hmac-sha1-96;
            encryption-algorithm 3des-cbc;
        }
        policy demo-policy {
            perfect-forward-secrecy {
                keys group2;
            }
            proposals ipsec-demo-proposal;
        }
        establish-tunnels immediately;
    }
    service-set service-set-dynamic-demo-service-set {
        next-hop-service {
            inside-service-interface ms-1/2/0.1;
            outside-service-interface ms-1/2/0.2;
        }
        ipsec-vpn-options {

```

```
        trusted-ca entrust;  
        local-gateway 10.1.15.2;  
    }  
    ipsec-vpn-rules rule-ike;  
} } }
```

Configuring Router 4

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI, at the [edit] hierarchy level, of Router 4.

```
set interfaces ge-0/0/0 description "to R3 ge-0/0/0"  
set interfaces ge-0/0/0 unit 0 family inet address 10.1.56.2/30  
set interfaces lo0 unit 0 family inet address 10.0.0.4/32  
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0  
set protocols ospf area 0.0.0.0 interface lo0.0  
set routing-options router-id 10.0.0.4
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To set up OSPF connectivity with Router 4

1. Configure the interfaces. In this step, you configure an Ethernet interface (ge-1/0/1) and the loopback interface.

```
[edit interfaces]  
user@router4# set ge-0/0/0 description "to R3 ge-0/0/0"  
user@router4# set ge-0/0/0 unit 0 family inet address 10.1.56.2/30  
user@router4# set lo0 unit 0 family inet address 10.0.0.4/32
```

2. Specify the OSPF area and associate the interfaces with the OSPF area.

```
[edit protocols]  
user@router4# set ospf area 0.0.0.0 interface ge-0/0/0  
user@router4# set ospf area 0.0.0.0 interface lo0.0
```

3. Configure the router ID.

```
[edit routing-options]  
user@router4# set router-id 10.0.0.4
```

Results From the configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration

```
user@router4# show interfaces  
interfaces {  
  ge-0/0/0 {  
    description "To R3 ge-0/0/0";
```

```

        unit 0 {
            family inet {
                address 10.1.56.2/30;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.0.0.4/32;
            }
        }
    }
}

user@router4# show protocols ospf
protocols {
    ospf {
        area 0.0.0.0 {
            interface ge-0/0/0.0;
            interface lo0.0;
        }
    }
}

user@router4# show routing-options
routing-options {
    router-id 10.0.0.4;
}

```

Verification

Verifying Your Work on Router 1

Purpose On Router 1, verify ping command to the so-0/0/0 interface on Router 4 to send traffic across the IPsec tunnel.

Action From operational mode, enter **ping 10.1.56.2**.

```

user@router1>ping 10.1.56.2
PING 10.1.56.2 (10.1.56.2): 56 data bytes
64 bytes from 10.1.56.2: icmp_seq=0 ttl=254 time=1.351 ms
64 bytes from 10.1.56.2: icmp_seq=1 ttl=254 time=1.187 ms
64 bytes from 10.1.56.2: icmp_seq=2 ttl=254 time=1.172 ms
64 bytes from 10.1.56.2: icmp_seq=3 ttl=254 time=1.154 ms
64 bytes from 10.1.56.2: icmp_seq=4 ttl=254 time=1.156 ms
^C
--- 10.1.56.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.154/1.204/1.351/0.074 ms

```

If you ping the loopback address of Router 4, the operation succeeds because the address is part of the OSPF network configured on Router 4.

```

user@router1>ping 10.0.0.4
PING 10.0.0.4 (10.0.0.4): 56 data bytes
64 bytes from 10.0.0.4: icmp_seq=0 ttl=62 time=1.318 ms
64 bytes from 10.0.0.4: icmp_seq=1 ttl=62 time=1.084 ms

```

```
64 bytes from 10.0.0.4: icmp_seq=2 ttl=62 time=3.260 ms
^C
--- 10.0.0.4 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.084/1.887/3.260/0.975 ms
```

Verifying Your Work on Router 2

Purpose To verify that matched traffic is being diverted to the bidirectional IPsec tunnel, view the IPsec statistics:

Action From operational mode, enter the **show services ipsec-vpn ipsec statistics**.

```
user@router2>show services ipsec-vpn ipsec statistics
PIC: sp-1/2/0, Service set: service-set-dynamic-demo-service-set
ESP Statistics:
Encrypted bytes: 162056
Decrypted bytes: 161896
Encrypted packets: 2215
Decrypted packets: 2216
AH Statistics:
Input bytes: 0
Output bytes: 0
Input packets: 0
Output packets: 0
Errors:
AH authentication failures: 0, Replay errors: 0
ESP authentication failures: 0, ESP decryption failures: 0
Bad headers: 0, Bad trailers: 0
```

To verify that the IKE SA negotiation is successful, issue the **show services ipsec-vpn ike security-associations** command:

From operational mode, enter the **show services ipsec-vpn ike security-associations**

```
user@router2> show services ipsec-vpn ike security-associations
Remote Address State Initiator cookie Responder cookie Exchange type
10.1.15.2 Matured d82610c59114fd37 ec4391f76783ef28 Main
```

To verify that the IPsec security association is active, issue the **show services ipsec-vpn ipsec security-associations detail** command. Notice that the SA contains the default settings inherent in the Services PIC, such as ESP for the protocol and HMAC-SHA1-96 for the authentication algorithm.

From operational mode, enter the **show services ipsec-vpn ipsec security-associations detail**

```
user@router2> show services ipsec-vpn ipsec security-associations detail
Service set: service-set-dynamic-demo-service-set
Rule: rule-ike, Term: term-ike, Tunnel index: 1
Local gateway: 10.1.15.1, Remote gateway: 10.1.15.2
IPsec inside interface: sp-1/2/0.1
Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Direction: inbound, SPI: 857451461, AUX-SPI: 0
Mode: tunnel, Type: dynamic, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Soft lifetime: Expires in 9052 seconds
Hard lifetime: Expires in 9187 seconds
```

```

Anti-replay service: Enabled, Replay window size: 64
Direction: outbound, SPI: 1272330309, AUX-SPI: 0
Mode: tunnel, Type: dynamic, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Soft lifetime: Expires in 9052 seconds
Hard lifetime: Expires in 9187 seconds
Anti-replay service: Enabled, Replay window size: 64

```

To display the digital certificates that are used to establish the IPsec tunnel, issue the `show services ipsec-vpn certificates` command:

From operational mode, enter the **show services ipsec-vpn certificates**

```

user@router2> show services ipsec-vpn certificates
Service set: service-set-dynamic-demo-service-set, Total entries: 3
Certificate cache entry: 3
Flags: Non-root Trusted
Issued to: router3.juniper.net, Issued by: juniper
Alternate subject: router3.juniper.net
Validity:
Not before: 2005 Nov 21st, 23:33:58 GMT
Not after: 2008 Nov 22nd, 00:03:58 GMT
Certificate cache entry: 2
Flags: Non-root Trusted
Issued to: router2.juniper.net, Issued by: juniper
Alternate subject: router2.juniper.net
Validity:
Not before: 2005 Nov 21st, 23:28:22 GMT
Not after: 2008 Nov 21st, 23:58:22 GMT
Certificate cache entry: 1
Flags: Root Trusted
Issued to: juniper, Issued by: juniper
Validity:
Not before: 2005 Oct 18th, 23:54:22 GMT
Not after: 2025 Oct 19th, 00:24:22 GMT

```

To display the CA certificate, issue the `show security pki ca-certificate detail` command. Notice that there are three separate certificates: one for certificate signing, one for key encipherment, and one for the CA's digital signature.

From operational mode, enter the **show security pki ca-certificate detail**

```

user@router2> show security pki ca-certificate detail
Certificate identifier: entrust
Certificate version: 3
Serial number: 4355 9235
Issuer:
Organization: juniper, Country: us
Subject:
Organization: juniper, Country: us
Validity:
Not before: 2005 Oct 18th, 23:54:22 GMT
Not after: 2025 Oct 19th, 00:24:22 GMT
Public key algorithm: rsaEncryption(1024 bits)
cb:9e:2d:c0:70:f8:ea:3c:f2:b5:f0:02:48:87:dc:68:99:a3:57:4f
0e:b9:98:0b:95:47:0d:1f:97:7c:53:17:dd:1a:f8:da:e5:08:d1:1c
78:68:1f:2f:72:9f:a2:cf:81:e3:ce:c5:56:89:ce:f0:97:93:fa:36
19:3e:18:7d:8c:9d:21:fe:1f:c3:87:8d:b3:5d:f3:03:66:9d:16:a7
bf:18:3f:f0:7a:80:f0:62:50:43:83:4f:0e:d7:c6:42:48:c0:8a:b2
c7:46:30:38:df:9b:dc:bc:b5:08:7a:f3:cd:64:db:2b:71:67:fe:d8
04:47:08:07:de:17:23:13

```

Signature algorithm: sha1WithRSAEncryption
Fingerprint:
00:8e:6f:58:dd:68:bf:25:0a:e3:f9:17:70:d6:61:f3:53:a7:79:10 (sha1)
71:6f:6a:76:17:9b:d6:2a:e7:5a:72:97:82:6d:26:86 (md5)
Distribution CRL:
C=us, O=juniper, CN=CRL1
http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: CRL signing, Certificate signing
Certificate identifier: entrust
Certificate version: 3
Serial number: 4355 925c
Issuer:
Organization: juniper, Country: us
Subject:
Organization: juniper, Country: us, Common name: First Officer
Validity:
Not before: 2005 Oct 18th, 23:55:59 GMT
Not after: 2008 Oct 19th, 00:25:59 GMT
Public key algorithm: rsaEncryption(1024 bits)
c0:a4:21:32:95:0a:cd:ec:12:03:d1:a2:89:71:8e:ce:4e:a6:f9:2f
1a:9a:13:8c:f6:a0:3d:c9:bd:9d:c2:a0:41:77:99:1b:1e:ed:5b:80
34:46:f8:5b:28:34:38:2e:91:7d:4e:ad:14:86:78:67:e7:02:1d:2e
19:11:b7:fa:0d:ba:64:20:e1:28:4e:3e:bb:6e:64:dc:cd:b1:b4:7a
ca:8f:47:dd:40:69:c2:35:95:ce:b8:85:56:d7:0f:2d:04:4d:5d:d8
42:e1:4f:6b:bf:38:c0:45:1e:9e:f0:b4:7f:74:6f:e9:70:fd:4a:78
da:eb:10:27:bd:46:34:33
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
bc:78:87:9b:a7:91:13:20:71:db:ac:b5:56:71:42:ad:1a:b6:46:17 (sha1)
23:79:40:c9:6d:a6:f0:ca:e0:13:30:d4:29:6f:86:79 (md5)
Distribution CRL:
C=us, O=juniper, CN=CRL1
http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: Key encipherment
Certificate identifier: entrust
Certificate version: 3
Serial number: 4355 925b
Issuer:
Organization: juniper, Country: us
Subject:
Organization: juniper, Country: us, Common name: First Officer
Validity:
Not before: 2005 Oct 18th, 23:55:59 GMT
Not after: 2008 Oct 19th, 00:25:59 GMT
Public key algorithm: rsaEncryption(1024 bits)
ea:75:c4:f3:58:08:ea:65:5c:7e:b3:de:63:0a:cf:cf:ec:9a:82:e2
d7:e8:b9:2f:bd:4b:cd:86:2f:f1:dd:d8:a2:95:af:ab:51:a5:49:4e
00:10:c6:25:ff:b5:49:6a:99:64:74:69:e5:8c:23:5b:b4:70:62:8e
e4:f9:a2:28:d4:54:e2:0b:1f:50:a2:92:cf:6c:8f:ae:10:d4:69:3c
90:e2:1f:04:ea:ac:05:9b:3a:93:74:d0:59:24:e9:d2:9d:c2:ef:22
b9:32:c7:2c:29:4f:91:cb:5a:26:fe:1d:c0:36:dc:f4:9c:8b:f5:26
af:44:bf:53:aa:d4:5f:67
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
46:71:15:34:f0:a6:41:76:65:81:33:4f:68:47:c4:df:78:b8:e3:3f (sha1)
ee:cc:c7:f4:5d:ac:65:33:0a:55:db:59:72:2c:dd:16 (md5)
Distribution CRL:
C=us, O=juniper, CN=CRL1
http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: Digital signature

To display the local certificate request, issue the `show security pki certificate-request` command:

From operational mode, enter the **show security pki certificate-request**

```
user@router2> show security pki certificate-request
Certificate identifier: local-entrust2
Issued to: router2.juniper.net
Public key algorithm: rsaEncryption(1024 bits)
Public key verification status: Passed
```

To display the local certificate, issue the `show security pki local-certificate` command:

From operational mode, enter the **show security pki local-certificate**

```
user@router2> show security pki local-certificate
Certificate identifier: local-entrust2
Issued to: router2.juniper.net, Issued by: juniper
Validity:
Not before: 2005 Nov 21st, 23:28:22 GMT
Not after: 2008 Nov 21st, 23:58:22 GMT
Public key algorithm: rsaEncryption(1024 bits)
Public key verification status: Passed
```

Verifying Your Work on Router 3

Purpose To verify that matched traffic is being diverted to the bidirectional IPsec tunnel, view the IPsec statistics:

Action From operational mode, enter the **show services ipsec-vpn ipsec statistics**.

```
user@router3>show services ipsec-vpn ipsec statistics
PIC: sp-1/2/0, Service set: service-set-dynamic-demo-service-set
ESP Statistics:
Encrypted bytes: 161896
Decrypted bytes: 162056
Encrypted packets: 2216
Decrypted packets: 2215
AH Statistics:
Input bytes: 0
Output bytes: 0
Input packets: 0
Output packets: 0
Errors:
AH authentication failures: 0, Replay errors: 0
ESP authentication failures: 0, ESP decryption failures: 0
Bad headers: 0, Bad trailers: 0
```

To verify that the IKE SA negotiation is successful, issue the `show services ipsec-vpn ike security-associations` command. To be successful, the SA on Router 3 must contain the same settings you specified on Router 2.

From operational mode, enter the **show services ipsec-vpn ike security-associations**.

```
user@router3>show services ipsec-vpn ike security-associations
Remote Address State Initiator cookie Responder cookie Exchange type
10.1.15.1 Matured d82610c59114fd37 ec4391f76783ef28 Main
```

To verify that the IPsec SA is active, issue the `show services ipsec-vpn ipsec security-associations detail` command. To be successful, the SA on Router 3 must contain the same settings you specified on Router 2.

From operational mode, enter the **show services ipsec-vpn ipsec security-associations detail**.

```
user@router3>show services ipsec-vpn ipsec security-associations detail
Service set: service-set-dynamic-demo-service-set
Rule: rule-ike, Term: term-ike, Tunnel index: 1
Local gateway: 10.1.15.2, Remote gateway: 10.1.15.1
IPsec inside interface: sp-1/2/0.1
Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Direction: inbound, SPI: 1272330309, AUX-SPI: 0
Mode: tunnel, Type: dynamic, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Soft lifetime: Expires in 7219 seconds
Hard lifetime: Expires in 7309 seconds
Anti-replay service: Enabled, Replay window size: 64
Direction: outbound, SPI: 857451461, AUX-SPI: 0
Mode: tunnel, Type: dynamic, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Soft lifetime: Expires in 7219 seconds
Hard lifetime: Expires in 7309 seconds
Anti-replay service: Enabled, Replay window size: 64
```

To display the digital certificates that are used to establish the IPsec tunnel, issue the `show services ipsec-vpn certificates` command:

From operational mode, enter the **show services ipsec-vpn certificates**.

```
user@router3>show services ipsec-vpn certificates
Service set: service-set-dynamic-demo-service-set, Total entries: 3
Certificate cache entry: 3
Flags: Non-root Trusted
Issued to: router3.juniper.net, Issued by: juniper
Alternate subject: router3.juniper.net
Validity:
Not before: 2005 Nov 21st, 23:33:58 GMT
Not after: 2008 Nov 22nd, 00:03:58 GMT
Certificate cache entry: 2
Flags: Non-root Trusted
Issued to: router2.juniper.net, Issued by: juniper
Alternate subject: router2.juniper.net
Validity:
Not before: 2005 Nov 21st, 23:28:22 GMT
Not after: 2008 Nov 21st, 23:58:22 GMT
Certificate cache entry: 1
Flags: Root Trusted
Issued to: juniper, Issued by: juniper
Validity:
Not before: 2005 Oct 18th, 23:54:22 GMT
Not after: 2025 Oct 19th, 00:24:22 GMT
```

To display the CA certificate, issue the `show security pki ca-certificate detail` command. Notice that there are three separate certificates: one for certificate signing, one for key encipherment, and one for the CA's digital signature.

From operational mode, enter the **show security pki ca-certificate detail**.

```

user@router3>show security pki ca-certificate detail
Certificate identifier: entrust
Certificate version: 3
Serial number: 4355 9235
Issuer:
Organization: juniper, Country: us
Subject:
Organization: juniper, Country: us
Validity:
Not before: 2005 Oct 18th, 23:54:22 GMT
Not after: 2025 Oct 19th, 00:24:22 GMT
Public key algorithm: rsaEncryption(1024 bits)
cb:9e:2d:c0:70:f8:ea:3c:f2:b5:f0:02:48:87:dc:68:99:a3:57:4f
0e:b9:98:0b:95:47:0d:1f:97:7c:53:17:dd:1a:f8:da:e5:08:d1:1c
78:68:1f:2f:72:9f:a2:cf:81:e3:ce:c5:56:89:ce:f0:97:93:fa:36
19:3e:18:7d:8c:9d:21:fe:1f:c3:87:8d:b3:5d:f3:03:66:9d:16:a7
bf:18:3f:f0:7a:80:f0:62:50:43:83:4f:0e:d7:c6:42:48:c0:8a:b2
c7:46:30:38:df:9b:dc:bc:b5:08:7a:f3:cd:64:db:2b:71:67:fe:d8
04:47:08:07:de:17:23:13
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
00:8e:6f:58:dd:68:bf:25:0a:e3:f9:17:70:d6:61:f3:53:a7:79:10 (sha1)
71:6f:6a:76:17:9b:d6:2a:e7:5a:72:97:82:6d:26:86 (md5)
Distribution CRL:
C=us, O=juniper, CN=CRL1
http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: CRL signing, Certificate signing
Certificate identifier: entrust
Certificate version: 3
Serial number: 4355 925c
Issuer:
Organization: juniper, Country: us
Subject:
Organization: juniper, Country: us, Common name: First Officer
Validity:
Not before: 2005 Oct 18th, 23:55:59 GMT
Not after: 2008 Oct 19th, 00:25:59 GMT
Public key algorithm: rsaEncryption(1024 bits)
c0:a4:21:32:95:0a:cd:ec:12:03:d1:a2:89:71:8e:ce:4e:a6:f9:2f
1a:9a:13:8c:f6:a0:3d:c9:bd:9d:c2:a0:41:77:99:1b:1e:ed:5b:80
34:46:f8:5b:28:34:38:2e:91:7d:4e:ad:14:86:78:67:e7:02:1d:2e
19:11:b7:fa:0d:ba:64:20:e1:28:4e:3e:bb:6e:64:dc:cd:b1:b4:7a
ca:8f:47:dd:40:69:c2:35:95:ce:b8:85:56:d7:0f:2d:04:4d:5d:d8
42:e1:4f:6b:bf:38:c0:45:1e:9e:f0:b4:7f:74:6f:e9:70:fd:4a:78
da:eb:10:27:bd:46:34:33
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
bc:78:87:9b:a7:91:13:20:71:db:ac:b5:56:71:42:ad:1a:b6:46:17 (sha1)
23:79:40:c9:6d:a6:f0:ca:e0:13:30:d4:29:6f:86:79 (md5)
Distribution CRL:
C=us, O=juniper, CN=CRL1
http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: Key encipherment
Certificate identifier: entrust
Certificate version: 3
Serial number: 4355 925b
Issuer:
Organization: juniper, Country: us
Subject:

```

```
Organization: juniper, Country: us, Common name: First Officer
Validity:
Not before: 2005 Oct 18th, 23:55:59 GMT
Not after: 2008 Oct 19th, 00:25:59 GMT
Public key algorithm: rsaEncryption(1024 bits)
ea:75:c4:f3:58:08:ea:65:5c:7e:b3:de:63:0a:cf:cf:ec:9a:82:e2
d7:e8:b9:2f:bd:4b:cd:86:2f:f1:dd:d8:a2:95:af:ab:51:a5:49:4e
00:10:c6:25:ff:b5:49:6a:99:64:74:69:e5:8c:23:5b:b4:70:62:8e
e4:f9:a2:28:d4:54:e2:0b:1f:50:a2:92:cf:6c:8f:ae:10:d4:69:3c
90:e2:1f:04:ea:ac:05:9b:3a:93:74:d0:59:24:e9:d2:9d:c2:ef:22
b9:32:c7:2c:29:4f:91:cb:5a:26:fe:1d:c0:36:dc:f4:9c:8b:f5:26
af:44:bf:53:aa:d4:5f:67
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
46:71:15:34:f0:a6:41:76:65:81:33:4f:68:47:c4:df:78:b8:e3:3f (sha1)
ee:cc:c7:f4:5d:ac:65:33:0a:55:db:59:72:2c:dd:16 (md5)
Distribution CRL:
C=us, O=juniper, CN=CRL1
http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: Digital signature
```

To display the local certificate request, issue the `show security pki certificate-request` command:

From operational mode, enter the **show security pki certificate-request**.

```
user@router3>show security pki certificate-request
Certificate identifier: local-entrust3
Issued to: router3.juniper.net
Public key algorithm: rsaEncryption(1024 bits)
Public key verification status: Passed
```

To display the local certificate, issue the `show security pki local-certificate` command:

From operational mode, enter the **show security pki local-certificate**.

```
user@router3>show security pki local-certificate
Certificate identifier: local-entrust3
Issued to: router3.juniper.net, Issued by: juniper
Validity:
Not before: 2005 Nov 21st, 23:33:58 GMT
Not after: 2008 Nov 22nd, 00:03:58 GMT
Public key algorithm: rsaEncryption(1024 bits)
Public key verification status: Passed
```

Verifying Your Work on Router 4

Purpose On Router 4, issue a ping command to the so-0/0/0 interface on Router 1 to send traffic across the IPsec tunnel.

Action From operational mode, enter **ping 10.1.12.2**.

```
user@router4>ping 10.1.12.2
PING 10.1.12.2 (10.1.12.2): 56 data bytes
64 bytes from 10.1.12.2: icmp_seq=0 ttl=254 time=1.350 ms
64 bytes from 10.1.12.2: icmp_seq=1 ttl=254 time=1.161 ms
64 bytes from 10.1.12.2: icmp_seq=2 ttl=254 time=1.124 ms
64 bytes from 10.1.12.2: icmp_seq=5 ttl=254 time=1.116 ms
^C
--- 10.1.12.2 ping statistics ---
```

```
4 packets transmitted, 4 packets received, 0% packet loss  
round-trip min/avg/max/stddev = 1.116/1.172/1.350/0.081 ms
```

The final way you can confirm that traffic travels over the IPsec tunnel is by issuing the `traceroute` command to the `so-0/0/0` interface on Router 1. Notice that the physical interface between Routers 2 and 3 is not referenced in the path; traffic enters the IPsec tunnel through the adaptive services IPsec inside interface on Router 3, passes through the loopback interface on Router 2, and ends at the `so-0/0/0` interface on Router 1.

From operational mode, enter the **`traceroute 10.1.12.2`**.

```
user@router4>traceroute 10.1.12.2  
traceroute to 10.1.12.2 (10.1.12.2), 30 hops max, 40 byte packets  
1 10.1.15.2 (10.1.15.2) 0.987 ms 0.630 ms 0.563 ms  
2 10.0.0.2 (10.0.0.2) 1.194 ms 1.058 ms 1.033 ms  
3 10.1.12.2 (10.1.12.2) 1.073 ms 0.949 ms 0.932 ms
```

**Related
Documentation**

- [Understanding Junos VPN Site Secure on page 371](#)
- [Configuring Security Associations on page 385](#)
- [Configuring IKE Proposals on page 405](#)
- [Configuring IKE Policies on page 409](#)
- [Example: Configuring IKE Dynamic SAs on page 466](#)
- [Example: Configuring Manual SAs on page 391](#)
- [Requesting for and Installing a Digital Certificates on Your Router](#)

Enabling IPsec for the Services SDK

- [Configuring Junos VPN Site Secure or IPSec VPN on page 507](#)

Configuring Junos VPN Site Secure or IPSec VPN

IPsec VPN is supported on all MX Series routers with MS-MICs, MS-MPCs, or MS-DPCs.

On M Series and T Series routers, IPsec VPN is supported with Multiservices 100 PICs, Multiservices 400 PICs, and Multiservices 500 PICs.

MS-MICs and MS-MPCs are supported from Junos OS Release 13.2 and later. MS-MICs and MS-MPCs support all features that are supported by MS-DPCs and MS-PICs except for authentication header protocol (ah), encapsulating security payload protocol (esp), and bundle (ah and esp protocol) protocol for a dynamic or manual security association and flowless IPsec service.

- Related Documentation**
- [Configuring Security Associations on page 385](#)
 - [Service Sets for IPsec Tunnels on page 430](#)

PART 8

Alleviating Congestion and Controlling Service Using CoS

- [Class of Service Overview on page 511](#)
- [Class of Service Configuration Overview on page 513](#)
- [Configuring Class of Service on LSQ Interfaces on page 523](#)

Class of Service Overview

- [Class of Service Overview on page 511](#)

Class of Service Overview

The CoS configuration available for the AS PIC enables you to configure Differentiated Services (DiffServ) code point (DSCP) marking and forwarding-class assignment for packets transiting the AS PIC. You can configure the CoS service alongside the stateful firewall and NAT services, using a similar rule structure. The component structures are described in detail in the *Class of Service Feature Guide for Routing Devices*.

Standards for Differentiated Services are described in the following documents:

- RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*
- RFC 2475, *An Architecture for Differentiated Services*



NOTE: CoS BA classification is not supported on services interfaces.

Related Documentation

- [Restrictions and Cautions for CoS Configuration on Services Interfaces on page 513](#)
- [Configuring CoS Rules on page 514](#)
- [Configuring CoS Rule Sets on page 519](#)
- [Examples: Configuring CoS on Services Interfaces on page 519](#)

Class of Service Configuration Overview

- [Restrictions and Cautions for CoS Configuration on Services Interfaces on page 513](#)
- [Configuring CoS Rules on page 514](#)
- [Configuring CoS Rule Sets on page 519](#)
- [Examples: Configuring CoS on Services Interfaces on page 519](#)

Restrictions and Cautions for CoS Configuration on Services Interfaces

The following restrictions and cautions apply to CoS configuration on services interfaces:

- The adaptive services interface does not support scheduling, only DiffServ marking and queue assignment. You must configure scheduling at the **[edit class-of-service]** hierarchy level on the output interface or fabric.
- In the default configuration, queues 1 and 2 receive 0 percent bandwidth. If packets will be assigned to these queues, you must configure a scheduling map.
- You must issue a **commit full** command before using custom forwarding-class names in the configuration.
- Only the Junos standard DiffServ names can be used in the configuration. Custom names are not recognized.
- On M Series routers, you can configure rewrite rules that change packet headers and attach the rules to output interfaces. These rules might overwrite the DSCP marking configured on an AS or MultiServices PIC. It is important to keep this adverse effect in mind and use care when creating system-wide configurations.

For example, knowing that the AS or MultiServices PIC can mark packets with any ToS or DSCP value and the output interface is restricted to only eight DSCP values, rewrite rules on the output interface condense the mapping from 64 to 8 values with overall loss of granularity. In this case, you have the following options:

- Remove the rewrite rules from the output interface.
- Configure the output interface to include the most important mappings.

Related Documentation

- [Class of Service Overview on page 511](#)
- [Configuring CoS Rules on page 514](#)

- [Configuring CoS Rule Sets on page 519](#)
- [Examples: Configuring CoS on Services Interfaces on page 519](#)

Configuring CoS Rules

To configure a CoS rule, include the **rule** *rule-name* statement at the **[edit services cos]** hierarchy level:

```
[edit services cos]
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application-sets set-name;
      applications [ application-names ];
      destination-address address;
      destination-prefix-list list-name <except>;
      source-address address;
      source-prefix-list list-name <except>;
    }
    then {
      application-profile profile-name;
      dscp (alias | bits);
      forwarding-class class-name;
      syslog;
      (reflexive | reverse) {
        application-profile profile-name;
        dscp (alias | bits);
        forwarding-class class-name;
        syslog;
      }
    }
  }
}
```

Each CoS rule consists of a set of terms, similar to a filter configured at the **[edit firewall]** hierarchy level. A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded.
- **then** statement—Specifies the actions and action modifiers to be performed by the router software.

The following sections explain how to configure the components of CoS rules:

- [Configuring Match Direction for CoS Rules on page 515](#)
- [Configuring Match Conditions In CoS Rules on page 515](#)
- [Configuring Actions in CoS Rules on page 516](#)
- [Example: Configuring CoS Rules on page 518](#)

Configuring Match Direction for CoS Rules

Each rule must include a **match-direction** statement that specifies the direction in which the rule match is applied. To configure where the match is applied, include the **match-direction** statement at the **[edit services cos rule rule-name]** hierarchy level:

```
match-direction (input | output | input-output);
```

If you configure **match-direction input-output**, bidirectional rule creation is allowed.

The match direction is used with respect to the traffic flow through the AS or Multiservices PIC. When a packet is sent to the PIC, direction information is carried along with it.

With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.

With a next-hop service set, packet direction is determined by the interface used to route the packet to the AS or Multiservices PIC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the AS or Multiservices PIC, the packet direction is output. For more information on inside and outside interfaces, see [“Configuring Service Sets to be Applied to Services Interfaces” on page 31](#).

On the AS or Multiservices PIC, a flow lookup is performed. If no flow is found, rule processing is performed. All rules in the service set are considered. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that matches the packet direction are considered.

Configuring Match Conditions In CoS Rules

To configure CoS match conditions, include the **from** statement at the **[edit services cos rule rule-name term term-name]** hierarchy level:

```
from {
  application-sets set-name;
  applications [ application-names ];
  destination-address address;
  destination-prefix-list list-name <except>;
  source-address address;
  source-prefix-list list-name <except>;
}
```

The source address and destination address can be either IPv4 or IPv6. You can use either the source address or the destination address as a match condition, in the same way that you would configure a firewall filter; for more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Alternatively, you can specify a list of source or destination prefixes by configuring the **prefix-list** statement at the **[edit policy-options]** hierarchy level and then including either the **destination-prefix-list** or **source-prefix-list** statement in the CoS rule. For an example, see [“Examples: Configuring Stateful Firewall Rules” on page 335](#).

If you omit the **from** term, the router accepts all traffic and the default protocol handlers take effect:

- User Datagram Protocol (UDP), Transmission Control Protocol (TCP), and Internet Control Message Protocol (ICMP) create a bidirectional flow with a predicted reverse flow.
- IP creates a unidirectional flow.

You can also include application protocol definitions you have configured at the **[edit applications]** hierarchy level; for more information, see [“Configuring Application Protocol Properties” on page 303](#).

- To apply one or more specific application protocol definitions, include the **applications** statement at the **[edit services cos rule rule-name term term-name from]** hierarchy level.
- To apply one or more sets of application protocol definitions you have defined, include the **application-sets** statement at the **[edit services cos rule rule-name term term-name from]** hierarchy level.



NOTE: If you include one of the statements that specifies application protocols, the router derives port and protocol information from the corresponding configuration at the **[edit applications]** hierarchy level; you cannot specify these properties as match conditions.

Configuring Actions in CoS Rules

To configure CoS actions, include the **then** statement at the **[edit services cos rule rule-name term term-name]** hierarchy level:

```
[edit services cos rule rule-name term term-name]
then {
  application-profile profile-name;
  dscp (alias | bits);
  forwarding-class class-name;
  syslog;
  (reflexive | reverse) {
    application-profile profile-name;
    dscp (alias | bits);
    forwarding-class class-name;
    syslog;
  }
}
```

The principal CoS actions are as follows:

- **dscp**—Causes the packet to be marked with the specified DiffServ code point (DSCP) value or alias.
- **forwarding-class**—Causes the packet to be assigned to the specified forwarding class.

For detailed information about DSCP values and forwarding classes, see “[Examples: Configuring CoS on Services Interfaces](#)” on page 519 or the *Class of Service Feature Guide for Routing Devices*.

You can optionally set the configuration to record information in the system logging facility by including the **syslog** statement at the **[edit services cos rule rule-name term term-name then]** hierarchy level. This statement overrides any **syslog** setting included in the service set or interface default configuration.

For information about some additional CoS actions, see the following sections:

- [Configuring Application Profiles for Use as CoS Rule Actions on page 517](#)
- [Configuring Reflexive and Reverse CoS Rule Actions on page 518](#)

Configuring Application Profiles for Use as CoS Rule Actions

You can optionally define one or more application profiles for inclusion in CoS actions. To configure application profiles, include the **application-profile** statement at the **[edit services cos]** hierarchy level:

```
[edit services cos]
application-profile profile-name {
  ftp {
    data {
      dscp (alias | bits);
      forwarding-class class-name;
    }
  }
  sip {
    video {
      dscp (alias | bits);
      forwarding-class class-name;
    }
    voice {
      dscp (alias | bits);
      forwarding-class class-name;
    }
  }
}
```

The **application-profile** statement includes two main components and three traffic types: **ftp** with the **data** traffic type and **sip** with the **video** and **voice** traffic types. You can set the appropriate **dscp** and **forwarding-class** values for each component within the application profile.



NOTE: The **ftp** and **sip** statements are not supported on Juniper Network MX Series 3D Universal Edge Routers.

You can apply the application profile to a CoS configuration by including it at the **[edit services cos rule rule-name term term-name then]** hierarchy level.

Configuring Reflexive and Reverse CoS Rule Actions

CoS services are unidirectional. It might be necessary to specify different treatments for flows in opposite directions.

Regardless of whether a packet matches the input, output or input-output direction, flows in both directions are created. A forward, reverse, or forward-and-reverse CoS action is associated with each flow. Bear in mind that the flow in the opposite direction might end up having a CoS action associated with it that you have not specifically configured.

To control the direction in which service is applied, as distinct from the direction in which the rule match is applied, you can configure the (**reflexive** | **reverse**) statement at the **[edit services cos rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services cos rule rule-name term term-name then]
(reflexive | reverse) {
  application-profile profile-name;
  dscp (alias | bits);
  forwarding-class class-name;
  syslog;
}
```

The two actions are mutually exclusive:

- **reflexive** causes the equivalent opposing CoS action to be applied to flows in the opposite direction.
- **reverse** allows you to define the CoS behavior for flows in the reverse direction.

If you omit the statement, data flows inherit the CoS behavior of the forward control flow.

Example: Configuring CoS Rules

The following example show a CoS configuration containing two rules, one for input matching on a specified application set and the other for output matching on a specified source address:

```
[edit services]
cos {
  rule my-cos-rule {
    match-direction input-output;
    term term1 {
      from {
        source-address 10.1.3.2/32;
        applications sip;
      }
      then {
        dscp ef;
        syslog;
      }
    }
    term term2 {
      from {
```



```

        destination-address 10.2.3.2;
        applications http;
    }
    then {
        dscp af21;
    }
}
}
}

```

**Related
Documentation**

- [Class of Service Overview on page 511](#)
- [Restrictions and Cautions for CoS Configuration on Services Interfaces on page 513](#)
- [Configuring CoS Rule Sets on page 519](#)
- [Examples: Configuring CoS on Services Interfaces on page 519](#)

Configuring CoS Rule Sets

The **rule-set** statement defines a collection of CoS rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then you specify the order of the rules by including the **rule-set** statement at the **[edit services cos]** hierarchy level with a **rule** statement for each rule:

```

rule-set rule-set-name {
    rule rule-name;
}

```

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

**Related
Documentation**

- [Class of Service Overview on page 511](#)
- [Restrictions and Cautions for CoS Configuration on Services Interfaces on page 513](#)
- [Configuring CoS Rules on page 514](#)
- [Examples: Configuring CoS on Services Interfaces on page 519](#)

Examples: Configuring CoS on Services Interfaces

To make settings consistent across Juniper Networks routers, you configure many CoS settings at the **[edit class-of-service]** hierarchy level to be used on services interfaces. When you commit this configuration along with what you configure at the **[edit services cos]** hierarchy level, these properties are applied to the AS or MultiServices PIC.

The following configuration examples at the **[edit class-of-service]** hierarchy level can be applied on services interfaces. For more information, see the *Class of Service Feature Guide for Routing Devices*.



NOTE: The first two configurations, mapping forwarding-class name to forwarding-class ID and mapping forwarding-class name to queue number, are mutually exclusive.

Mapping Forwarding-Class Name to Forwarding-Class ID	<p>Map forwarding-class names to forwarding-class IDs:</p> <pre>[edit class-of-service] forwarding-classes { forwarding-class fc0 0; forwarding-class fc1 0; forwarding-class fc2 1; forwarding-class fc3 1; forwarding-class fc4 2; forwarding-class fc5 2; forwarding-class fc6 3; forwarding-class fc7 3; forwarding-class fc8 4; forwarding-class fc9 4; forwarding-class fc10 5; forwarding-class fc11 5; forwarding-class fc12 6; forwarding-class fc13 6; forwarding-class fc14 7; forwarding-class fc15 7; }</pre>
Mapping Forwarding-Class Name to Queue Number	<p>Map forwarding-class names to queue numbers:</p> <pre>[edit class-of-service] forwarding-classes { queue 0 be; queue 1 ef; queue 2 af; queue 3 nc; queue 4 ef1; queue 5 ef2; queue 6 af1; queue 7 nc1; }</pre>
Mapping Diffserv Code Point Aliases to DSCP Bits	<p>Map alias names to DSCP bit values. The aliases then can be used instead of the DSCP bits in adaptive services configurations.</p> <pre>[edit class-of-service] code-point-aliases { (dscp dscp-ipv6 exp ieee-802.1 inet-precedence) { alias bits; } }</pre>

Here is an example:

```
code-point-aliases {  
  dscp {  
    my1 110001;  
    my2 101110;  
    be 000001;  
    cs7 110000;  
  }  
}
```

**Related
Documentation**

- [Class of Service Overview on page 511](#)
- [Restrictions and Cautions for CoS Configuration on Services Interfaces on page 513](#)
- [Configuring CoS Rules on page 514](#)
- [Configuring CoS Rule Sets on page 519](#)

Configuring Class of Service on LSQ Interfaces

- [Configuring CoS Scheduling Queues on Logical LSQ Interfaces on page 523](#)
- [Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 527](#)
- [Configuring Link Services and CoS on Services PICs on page 529](#)
- [Oversubscribing Interface Bandwidth on LSQ Interfaces on page 532](#)
- [Configuring Guaranteed Minimum Rate on LSQ Interfaces on page 537](#)

Configuring CoS Scheduling Queues on Logical LSQ Interfaces

For link services IQ (**lsq-**) interfaces, you can specify a scheduler map for each logical unit. A logical unit represents either an MLPPP bundle or a DLCI configured on a FRF.16 bundle. The scheduler is applied to the traffic sent to an AS or Multiservices PIC running the Layer 2 link services package.

If you configure a scheduler map on a bundle, you must include the **per-unit-scheduler** statement at the **[edit interfaces lsq-fpc/pic/port]** hierarchy level. If you configure a scheduler map on an FRF.16 DLCI, you must include the **per-unit-scheduler** statement at the **[edit interfaces lsq-fpc/pic/port:channel]** hierarchy level. For more information, see the *Class of Service Feature Guide for Routing Devices*.

If you need latency guarantees for multiclass or LFI traffic, you must use channelized IQ PICs for the constituent links. With non-IQ PICs, because queueing is not done at the channelized interface level on the constituent links, latency-sensitive traffic might not receive the type of service that it should. Constituent links from the following PICs support latency guarantees:

- Channelized E1 IQ PIC
- Channelized OC3 IQ PIC
- Channelized OC12 IQ PIC
- Channelized STM1 IQ PIC
- Channelized T3 IQ PIC

For scheduling queues on a logical interface, you can configure the following scheduler map properties at the **[edit class-of-service schedulers]** hierarchy level:

- **buffer-size**—The queue size; for more information, see [“Configuring Scheduler Buffer Size” on page 524](#).
- **priority**—The transmit priority (low, high, strict-high); for more information, see [“Configuring Scheduler Priority” on page 525](#).
- **shaping-rate**—The subscribed transmit rate; for more information, see [“Configuring Scheduler Shaping Rate” on page 525](#).
- **drop-profile-map**—The random early detection (RED) drop profile; for more information, see [“Configuring Drop Profiles” on page 525](#).

When you configure MLPPP and FRF.12 on M Series and T Series routers, you should configure a single scheduler with non-zero percent transmission rates and buffer sizes for queues 0 through 3, and assign this scheduler to the link services IQ interface (**lsq**) and to each constituent link.

When you configure FRF.16 on M Series and T Series routers, you can assign a single scheduler map to the link services IQ interface (**lsq**) and to each link services IQ DLCI, or you can assign different scheduler maps to the various DLCIs of the bundle, as shown in [“Example: Configuring an LSQ Interface as an NxT1 Bundle Using FRF.16” on page 574](#). For the constituent links of an FRF.16 bundle, you do not need to configure a custom scheduler. Because LFI and multiclass are not supported for FRF.16, the traffic from each constituent link is transmitted from queue 0. This means you should allow most of the bandwidth to be used by queue 0. The default scheduler transmission rate and buffer size percentages for queues 0 through 3 are 95, 0, 0, and 5 percent, respectively. This default scheduler sends all user traffic to queue 0 and all network-control traffic to queue 3, and therefore it is well suited to the behavior of FRF.16. You can configure a custom scheduler that explicitly replicates the 95, 0, 0, and 5 percent queuing behaviors, and apply it to the constituent links.



NOTE: On T Series and M320 routers, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

For link services IQ interfaces (**lsq**), these scheduling properties work as they do in other PICs, except as noted in the following sections.



NOTE: On T Series and M320 routers, **lsq** interfaces do not support DiffServ code point (DSCP) and DSCP-IPv6 rewrite markers.

Configuring Scheduler Buffer Size

You can configure the scheduler buffer size in three ways: as a temporal value, as a percentage, and as a remainder. On a single logical interface (MLPPP or a FRF.16 DLCI), each queue can have a different buffer size.

If you specify a temporal value, the queuing algorithm starts dropping packets when it queues more than a computed number of bytes. This number is computed by multiplying logical interface speed by the temporal value. For MLPPP bundles, logical interface speed is equal to the bundle bandwidth, which is the sum of constituent link speeds minus link-layer overhead. For MLFR FRF.16 DLCIs, logical interface speed is equal to bundle bandwidth multiplied by the DLCI shaping rate. In all cases, the maximum temporal value is limited to 200 milliseconds.

Buffer size percentages are implicitly converted into temporal values by multiplying the percentage by 200 milliseconds. For example, buffer size specified as **buffer-size percent 20** is the same as a 40-millisecond temporal delay. The link services IQ implementation guarantees 200 milliseconds of buffer delay for all interfaces with T1 and higher speeds. For slower interfaces, it guarantees one second of buffer delay.

The queueing algorithm evenly distributes leftover bandwidth among all queues that are configured with the **buffer-size remainder** statement. The queueing algorithm guarantees enough space in the transmit buffer for two MTU-sized packets.

Configuring Scheduler Priority

The transmit priority of each queue is determined by the scheduler and the forwarding class. Each queue receives a guaranteed amount of bandwidth specified with the scheduler **transmit-rate** statement.

Configuring Scheduler Shaping Rate

You use the shaping rate to set the percentage of total bundle bandwidth that is dedicated to a DLCI. For link services IQ DLCIs, only percentages are accepted, which allows adjustments in response to dynamic changes in bundle bandwidth—for example, when a link goes up or down. This means that absolute shaping rates are not supported on FRF.16 bundles. Absolute shaping rates are allowed for MLPPP and MLFR bundles only.

For scheduling between DLCIs in a MLFR FRF.16 bundle, you can configure a shaping rate for each DLCI. A shaping rate is expressed as a percentage of the aggregate bundle bandwidth. Shaping rate percentages for all DLCIs within a bundle can add up to 100 percent or less. Leftover bandwidth is distributed equally to DLCIs that do not have the **shaping-rate** statement included at the **[edit class-of-service interfaces lsq-fpc/pic/port:channel unit logical-unit-number]** hierarchy level. If none of the DLCIs in an MLFR FRF.16 bundle specify a DLCI scheduler, the total bandwidth is evenly divided across all DLCIs.



NOTE: For FRF.16 bundles on link services IQ interfaces, only shaping rates based on percentage are supported.

Configuring Drop Profiles

You can configure random early detection (RED) on LSQ interfaces as in other CoS scenarios. To configure RED, include one or more drop profiles and attach them to a scheduler for a particular forwarding class. For more information about RED profiles, see the *Class of Service Feature Guide for Routing Devices*.

The LSQ implementation performs tail RED. It supports a maximum of 256 drop profiles per PIC. Drop profiles are configurable on a per-queue, per-loss-priority, and per-TCP-bit basis.

You can attach scheduler maps with configured RED drop profiles to any LSQ logical interface: an MLPPP bundle, an FRF.15 bundle, or an FRF.16 DLCI. Different queues (forwarding classes) on the same logical interface can have different associated drop profiles.

The following example shows how to configure a RED profile on an LSQ interface:

```
[edit]
class-of-service {
  drop-profiles {
    drop-low {
      # Configure suitable drop profile for low loss priority
      ...
    }
    drop-high {
      # Configure suitable drop profile for high loss priority
      ...
    }
  }
  scheduler-maps {
    schedmap {
      # Best-effort queue will use be-scheduler
      # Other queues may use different schedulers
      forwarding-class be scheduler be-scheduler;
      ...
    }
  }
  schedulers {
    be-scheduler {
      # Configure two drop profiles for low and high loss priority
      drop-profile-map loss-priority low protocol any drop-profile drop-low;
      drop-profile-map loss-priority high protocol any drop-profile drop-high;
      # Other scheduler parameters (buffer-size, priority,
      # and transmit-rate) are already supported.
      ...
    }
  }
  interfaces {
    lsq-1/3/0.0 {
      # Attach a scheduler map (that includes RED drop profiles)
      # to a LSQ logical interface.
      scheduler-map schedmap;
    }
  }
}
```



NOTE: The RED profiles should be applied only on the LSQ bundles and not on the egress links that constitute the bundle.

- Related Documentation**
- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
 - [Configuring Link Services and CoS on Services PICs on page 529](#)
 - [Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 527](#)
 - [Link Services Configuration for Junos Interfaces](#)

Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces

For link services IQ (**lsq-**) interfaces, you can specify fragmentation properties for specific forwarding classes. Traffic on each forwarding class can be either multilink encapsulated (fragmented and sequenced) or nonencapsulated (hashed with no fragmentation). By default, traffic in all forwarding classes is multilink encapsulated.

When you do not configure fragmentation properties for the queues on MLPPP interfaces, the fragmentation threshold you set at the **[edit interfaces *interface-name* unit *logical-unit-number* fragment-threshold]** hierarchy level is the fragmentation threshold for all forwarding classes within the MLPPP interface. For MLFR FRF.16 interfaces, the fragmentation threshold you set at the **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options fragment-threshold]** hierarchy level is the fragmentation threshold for all forwarding classes within the MLFR FRF.16 interface.

If you do not set a maximum fragment size anywhere in the configuration, packets are still fragmented if they exceed the smallest maximum transmission unit (MTU) or maximum received reconstructed unit (MRRU) of all the links in the bundle. A nonencapsulated flow uses only one link. If the flow exceeds a single link, then the forwarding class must be multilink encapsulated, unless the packet size exceeds the MTU/MRRU.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the MRRU by including the **mrru** statement at the **[edit interfaces *lsq-fpc/pic/port* unit *logical-unit-number*]** or **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]** hierarchy level. The MRRU is similar to the MTU, but is specific to link services interfaces. By default the MRRU size is 1500 bytes, and you can configure it to be from 1500 through 4500 bytes. For more information, see [“Configuring MRRU on Multilink and Link Services Logical Interfaces” on page 734](#).

To configure fragmentation properties on a queue, include the **fragmentation-maps** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      (fragment-threshold bytes | no-fragmentation);
      multilink-class number;
    }
  }
}
```

To set a per-forwarding class fragmentation threshold, include the **fragment-threshold** statement in the fragmentation map. This statement sets the maximum size of each multilink fragment.

To set traffic on a queue to be nonencapsulated rather than multilink encapsulated, include the **no-fragmentation** statement in the fragmentation map. This statement specifies that an extra fragmentation header is not prepended to the packets received on this queue and that static link load balancing is used to ensure in-order packet delivery.

For a given forwarding class, you can include either the **fragment-threshold** or **no-fragmentation** statement; they are mutually exclusive.

You use the **multilink-class** statement to map a forwarding class into a multiclass MLPPP (MCML). For a given forwarding class, you can include either the **multilink-class** or **no-fragmentation** statement; they are mutually exclusive. For more information about MCML, see [“Configuring Multiclass MLPPP on LSQ Interfaces” on page 562](#).

To associate a fragmentation map with a multilink PPP interface or MLFR FRF.16 DLCI, include the **fragmentation-map** statement at the **[edit class-of-service interfaces interface-name unit logical-unit-number]** hierarchy level:

```
[edit class-of-service interfaces]
lsq-fpc/pic/port {
  unit logical-unit-number { # Multilink PPP
    fragmentation-map map-name;
  }
lsq-fpc/pic/port:channel { # MLFR FRF.16
  unit logical-unit-number {
    fragmentation-map map-name;
  }
}
```

For configuration examples, see the following topics:

- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using MLPPP on page 565](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.16 on page 571](#)
- [Configuring LSQ Interfaces for Single Fractional T1 or El Interfaces Using MLPPP and LFI on page 577](#)
- [Configuring LSQ Interfaces for Single Fractional T1 or El Interfaces Using FRF.12 on page 582](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.15 on page 576](#)
- [Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 589](#)
- [Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12 on page 591](#)
- [Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP on page 593](#)

For Link Services PIC link services (**ls-**) interfaces, fragmentation maps are not supported. Instead, you enable LFI by including the **interleave-fragments** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level. For more information,

see [“Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces” on page 773](#).

Related Documentation

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Configuring Link Services and CoS on Services PICs on page 529](#)
- [Configuring CoS Scheduling Queues on Logical LSQ Interfaces on page 523](#)
- [Link Services Configuration for Junos Interfaces](#)

Configuring Link Services and CoS on Services PICs

To configure link services and CoS on an AS or Multiservices PIC, you must perform the following steps:

1. Enable the Layer 2 service package. You enable service packages per PIC, not per port. When you enable the Layer 2 service package, the entire PIC uses the configured package. To enable the Layer 2 service package, include the **service-package** statement at the **[edit chassis fpc slot-number pic pic-number adaptive-services]** hierarchy level, and specify **layer-2**:

```
[edit chassis fpc slot-number pic pic-number adaptive-services]
service-package layer-2;
```

For more information about AS or Multiservices PIC service packages, see [“Enabling Service Packages” on page 11](#) and [“Layer 2 Service Package Capabilities and Interfaces” on page 543](#).

2. Configure a multilink PPP or FRF.16 bundle by combining constituent links into a virtual link, or bundle.

Configuring an MLPPP Bundle

To configure an MLPPP bundle, configure constituent links and bundle properties by including the following statements in the configuration:

```
[edit interfaces interface-name unit logical-unit-number]
encapsulation ppp;
family mlppp {
    bundle lsq-fpc/pic/port.logical-unit-number;
}
[edit interfaces lsq-fpc/pic/port unit logical-unit-number]
drop-timeout milliseconds;
encapsulation multilink-ppp;
fragment-threshold bytes;
link-layer-overhead percent;
minimum-links number;
mrru bytes;
short-sequence;
family inet {
    address address;
}
```

For more information about these statements, see the *Link and Multilink Services Interfaces Feature Guide for Routing Devices*.

Configuring an MLFR FRF.16 Bundle

To configure an MLFR FRF.16 bundle, configure constituent links and bundle properties by including the following statements in the configuration:

```
[edit chassis fpc slot-number pic slot-number]
mlfr-uni-nni-bundles number;
[edit interfaces interface-name ]
encapsulation multilink-frame-relay-uni-nni;
unit logical-unit-number {
    family mlfr-uni-nni {
        bundle lsq-fpc/pic/port:channel;
    }
}
```

For more information about the **mlfr-uni-nni-bundles** statement, see the *Junos OS Administration Library for Routing Devices*. MLFR FRF.16 uses channels as logical units.

For MLFR FRF.16, you must configure one end as data circuit-terminating equipment (DCE) by including the following statements at the **[edit interfaces lsq-fpc/pic/port:channel]** hierarchy level.

```
encapsulation multilink-frame-relay-uni-nni;
dce;
mlfr-uni-nni-options {
    acknowledge-retries number;
    acknowledge-timer milliseconds;
    action-red-differential-delay (disable-tx | remove-link);
    drop-timeout milliseconds;
    fragment-threshold bytes;
    hello-timer milliseconds;
    link-layer-overhead percent;
    lmi-type (ansi | itu);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;
    red-differential-delay milliseconds;
    t391 number;
    t392 number;
    yellow-differential-delay milliseconds;
}
unit logical-unit-number {
    dlci dlci-identifier;
    family inet {
        address address;
    }
}
```

For more information about MLFR UNI NNI properties, see *Link and Multilink Services Interfaces Feature Guide for Routing Devices*.

3. To configure CoS components for each multilink bundle, enable per-unit scheduling on the interface, configure a scheduler map, apply the scheduler to each queue, configure a fragmentation map, and apply the fragmentation map to each bundle. Include the following statements:

```

[edit interfaces]
lsq-fpc/pic/port {
  per-unit-scheduler; # Enables per-unit scheduling on the bundle
}
[edit class-of-service]
interfaces {
  lsq-fpc/pic/port { # Multilink PPP
    unit logical-unit-number {
      scheduler-map map-name; # Applies scheduler map to each queue
    }
  }
  lsq-fpc/pic/port:channel { # MLFR FRF.16
    unit logical-unit-number {
      # Scheduler map provides scheduling information for
      # the queues within a single DLCI.
      scheduler-map map-name;
      shaping-rate percent percent;
    }
  }
  forwarding-classes {
    queue queue-number class-name priority (high | low);
  }
  scheduler-maps {
    map-name {
      forwarding-class class-name scheduler scheduler-name;
    }
  }
  schedulers {
    scheduler-name {
      buffer-size (percent percentage | remainder | temporal microseconds);
      priority priority-level;
      transmit-rate (percent percentage | rate | remainder) <exact>;
    }
  }
  fragmentation-maps {
    map-name {
      forwarding-class class-name {
        fragment-threshold bytes;
        no-fragmentation;
      }
    }
  }
}

```

Associate a fragmentation map with a multilink PPP interface or MLFR FRF.16 DLCI by including the following statements at the **[edit class-of-service]** hierarchy level:

```

interfaces {
  lsq-fpc/pic/port {
    unit logical-unit-number { # Multilink PPP
      fragmentation-map map-name;
    }
  }
  lsq-fpc/pic/port:channel { # MLFR FRF.16
    unit logical-unit-number {
      fragmentation-map map-name;
    }
  }
}

```

Related Documentation

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS on page 545](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using SONET APS on page 548](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 548](#)
- [Link Services Configuration for Junos Interfaces](#)

Oversubscribing Interface Bandwidth on LSQ Interfaces

The term *oversubscribing interface bandwidth* means configuring shaping rates (peak information rates [PIRs]) so that their sum exceeds the interface bandwidth.

On Channelized IQ PICs, Gigabit Ethernet IQ PICs, and FRF.16 link services IQ (**lsq-**) interfaces on AS and Multiservices PICs, you can oversubscribe interface bandwidth. The logical interfaces (and DLCIs within an FRF.16 bundle) can be oversubscribed when there is leftover bandwidth. The oversubscription is limited to the configured PIR. Any unused bandwidth is distributed equally among oversubscribed logical interfaces or DLCIs.

For networks that are not likely to experience congestion, oversubscribing interface bandwidth improves network utilization, thereby allowing more customers to be provisioned on a single interface. If the actual data traffic does not exceed the interface bandwidth, oversubscription allows you to sell more bandwidth than the interface can support.

We recommend avoiding oversubscription in networks that are likely to experience congestion. Be careful not to oversubscribe a service by too much, because this can cause degradation in the performance of the router during congestion. When you configure oversubscription, some output queues can be starved if the actual data traffic exceeds the physical interface bandwidth. You can prevent degradation by using statistical multiplexing to ensure that the actual data traffic does not exceed the interface bandwidth.



NOTE: You cannot oversubscribe interface bandwidth when you configure traffic shaping using the method described in *Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs*.

When configuring oversubscription for FRF.16 bundle interfaces, you can assign traffic control profiles that apply on a physical interface basis. When you apply traffic control profiles to FRF.16 bundles at the *logical* interface level, member link interface bandwidth is underutilized when there is a small proportion of traffic or no traffic at all on an individual DLCI. Support for traffic control features on the FRF.16 bundle physical interface level addresses this limitation.

To configure oversubscription of an interface, perform the following steps:

1. Include the **shaping-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

```
[edit class-of-service traffic-control-profiles profile-name]
shaping-rate (percent percentage | rate);
```



NOTE: When configuring oversubscription for FRF.16 bundle interfaces on a physical interface basis, you *must* specify **shaping-rate** as a percentage.

On LSQ interfaces, you can configure the shaping rate as a percentage.

On IQ and IQ2 interfaces, you can configure the shaping rate as an absolute rate from 1000 through 6,400,000,000,000 bits per second.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the **shaping-rate** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level. However, with this configuration approach, you cannot independently control the delay-buffer rate, as described in Step 2.



NOTE: For channelized and Gigabit Ethernet IQ interfaces, the **shaping-rate** and **guaranteed-rate** statements are mutually exclusive. You cannot configure some logical interfaces to use a shaping rate and others to use a guaranteed rate. This means there are no service guarantees when you configure a PIR. For these interfaces, you can configure either a PIR or a committed information rate (CIR), but not both.

This restriction does not apply to Gigabit Ethernet IQ2 PICs or link services IQ (LSQ) interfaces on AS or Multiservices PICs. For LSQ and Gigabit Ethernet IQ2 interfaces, you can configure both a PIR and a CIR on an interface. For more information about CIRs, see [“Configuring Guaranteed Minimum Rate on LSQ Interfaces” on page 537](#).

2. Optionally, you can base the delay buffer calculation on a delay-buffer rate. To do this, include the **delay-buffer-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:



NOTE: When configuring oversubscription for FRF.16 bundle interfaces on a physical interface basis, you *must* specify **delay-buffer-rate** as a percentage.

```
[edit class-of-service traffic-control-profiles profile-name]
delay-buffer-rate (percent percentage | rate);
```

The delay-buffer rate overrides the shaping rate as the basis for the delay-buffer calculation. In other words, the shaping rate or scaled shaping rate is used for delay-buffer calculations only when the delay-buffer rate is not configured.

For LSQ interfaces, if you do not configure a delay-buffer rate, the guaranteed rate (CIR) is used to assign buffers. If you do not configure a guaranteed rate, the shaping rate (PIR) is used in the undersubscribed case, and the scaled shaping rate is used in the oversubscribed case.

On LSQ interfaces, you can configure the delay-buffer rate as a percentage.

On IQ and IQ2 interfaces, you can configure the delay-buffer rate as an absolute rate from 1000 through 6,400,000,000,000 bits per second.

The actual delay buffer is based on the calculations described in the *Class of Service Feature Guide for Routing Devices*. For an example showing how the delay-buffer rates are applied, see [“Examples: Oversubscribing an LSQ Interface” on page 535](#).

Configuring large buffers on relatively low-speed links can cause packet aging. To help prevent this problem, the software requires that the sum of the delay-buffer rates be less than or equal to the port speed.

This restriction does not eliminate the possibility of packet aging, so you should be cautious when using the **delay-buffer-rate** statement. Though some amount of extra buffering might be desirable for burst absorption, delay-buffer rates should not far exceed the service rate of the logical interface.

If you configure delay-buffer rates so that the sum exceeds the port speed, the configured delay-buffer rate is not implemented for the last logical interface that you configure. Instead, that logical interface receives a delay-buffer rate of zero, and a warning message is displayed in the CLI. If bandwidth becomes available (because another logical interface is deleted or deactivated, or the port speed is increased), the configured delay-buffer-rate is reevaluated and implemented if possible.

If you do not configure a delay-buffer rate or a guaranteed rate, the logical interface receives a delay-buffer rate in proportion to the shaping rate and the remaining delay-buffer rate available. In other words, the delay-buffer rate for each logical interface with no configured delay-buffer rate is equal to:

$$(\text{remaining delay-buffer rate} * \text{shaping rate}) / (\text{sum of shaping rates})$$

The remaining delay-buffer rate is equal to:

$$(\text{interface speed}) - (\text{sum of configured delay-buffer rates})$$

3. To assign a scheduler map to the logical interface, include the **scheduler-map** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

```
[edit class-of-service traffic-control-profiles profile-name]  
scheduler-map map-name;
```

For information about configuring schedulers and scheduler maps, see the *Class of Service Feature Guide for Routing Devices*.

4. Optionally, you can enable large buffer sizes to be configured. To do this, include the **q-pic-large-buffer** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]  
q-pic-large-buffer;
```


If you do not include this statement, the delay-buffer size is more restricted.

We recommend restricted buffers for delay-sensitive traffic, such as voice traffic. For more information, see the *Class of Service Feature Guide for Routing Devices*.

5. To enable scheduling on logical interfaces, include the **per-unit-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name ]
  per-unit-scheduler;
```

When you include this statement, the maximum number of VLANs supported is 768 on a single-port Gigabit Ethernet IQ PIC. On a two-port Gigabit Ethernet IQ PIC, the maximum number is 384.

6. To enable scheduling for FRF.16 bundles physical interfaces, include the **no-per-unit-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
  no-per-unit-scheduler;
```

7. To apply the traffic-scheduling profile to the logical interface, include the **output-traffic-control-profile** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
  output-traffic-control-profile profile-name;
```

You cannot include the **output-traffic-control-profile** statement in the configuration if any of the following statements are included in the logical interface configuration: **scheduler-map**, **shaping-rate**, **adaptive-shaper**, or **virtual-channel-group**.

For a table that shows how the bandwidth and delay buffer are allocated in various configurations, see the *Class of Service Feature Guide for Routing Devices*.

Examples: Oversubscribing an LSQ Interface

Oversubscribing an LSQ Interface with Scheduling Based on the Logical Interface

Apply a traffic-control profile to a logical interface representing a DLCI on an FRF.16 bundle.

```
interfaces {
  lsq-1/3/0:0 {
    per-unit-scheduler;
    unit 0 {
      dlci 100;
    }
    unit 1 {
      dlci 200;
    }
  }
}
class-of-service {
  traffic-control-profiles {
    tc_0 {
      shaping-rate percent 100;
      guaranteed-rate percent 60;
      delay-buffer-rate percent 80;
    }
  }
}
```

```
tc_1 {
    shaping-rate percent 80;
    guaranteed-rate percent 40;
}
}
interfaces {
    lsq-1/3/0 {
        unit 0 {
            output-traffic-control-profile tc_0;
        }
        unit 1 {
            output-traffic-control-profile tc_1;
        }
    }
}
}
```

**Oversubscribing an
LSQ Interface with
Scheduling Based on
the Physical Interface**

Apply a traffic-control profile to the physical interface representing an FRF.16 bundle:

```
interfaces {
    lsq-0/2/0:0 {
        no-per-unit-scheduler;
        encapsulation multilink-frame-relay-uni-nni;
        unit 0 {
            dlc1 100;
            family inet {
                address 18.18.18.2/24;
            }
        }
    }
}
class-of-service {
    traffic-control-profiles {
        rlsq_tc {
            scheduler-map rlsq;
            shaping-rate percent 60;
            delay-buffer-rate percent 10;
        }
    }
    interfaces {
        lsq-0/2/0:0 {
            output-traffic-control-profile rlsq_tc;
        }
    }
}
scheduler-maps {
    rlsq {
        forwarding-class best-effort scheduler rlsq_scheduler;
        forwarding-class expedited-forwarding scheduler rlsq_scheduler1;
    }
}
schedulers {
    rlsq_scheduler {
        transmit-rate percent 20;
        priority low;
    }
    rlsq_scheduler1 {
```

```

        transmit-rate percent 40;
        priority high;
    }
}

```

**Related
Documentation**

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces on page 561](#)
- [Configuring Guaranteed Minimum Rate on LSQ Interfaces on page 537](#)
- *Link Services Configuration for Junos Interfaces*

Configuring Guaranteed Minimum Rate on LSQ Interfaces

On Gigabit Ethernet IQ PICs, Channelized IQ PICs, and FRF.16 link services IQ (LSQ) interfaces on AS and Multiservices PICs, you can configure guaranteed bandwidth, also known as a committed information rate (CIR). This allows you to specify a guaranteed rate for each logical interface. The guaranteed rate is a minimum. If excess physical interface bandwidth is available for use, the logical interface receives more than the guaranteed rate provisioned for the interface.

You cannot provision the sum of the guaranteed rates to be more than the physical interface bandwidth, or the bundle bandwidth for LSQ interfaces. If the sum of the guaranteed rates exceeds the interface or bundle bandwidth, the commit operation does not fail, but the software automatically decreases the rates so that the sum of the guaranteed rates is equal to the available bundle bandwidth.

To configure a guaranteed minimum rate, perform the following steps:

1. Include the **guaranteed-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

```

[edit class-of-service traffic-control-profiles profile-name]
  guaranteed-rate (percent percentage | rate);

```

On LSQ interfaces, you can configure the guaranteed rate as a percentage.

On IQ and IQ2 interfaces, you can configure the guaranteed rate as an absolute rate from 1000 through 160,000,000,000 bits per second.



NOTE: For channelized and Gigabit Ethernet IQ interfaces, the **shaping-rate** and **guaranteed-rate** statements are mutually exclusive. You cannot configure some logical interfaces to use a shaping rate and others to use a guaranteed rate. This means there are no service guarantees when you configure a PIR. For these interfaces, you can configure either a PIR or a committed information rate (CIR), but not both.

This restriction does not apply to Gigabit Ethernet IQ2 PICs or link services IQ (LSQ) interfaces on AS or Multiservices PICs. For LSQ and Gigabit Ethernet IQ2 interfaces, you can configure both a PIR and a CIR on an interface. For more information about CIRs, see the *Class of Service Feature Guide for Routing Devices*.

2. Optionally, you can base the delay buffer calculation on a delay-buffer rate. To do this, include the **delay-buffer-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

```
[edit class-of-service traffic-control-profiles profile-name]  
delay-buffer-rate (percent percentage | rate);
```

On LSQ interfaces, you can configure the delay-buffer rate as a percentage.

On IQ and IQ2 interfaces, you can configure the delay-buffer rate as an absolute rate from 1000 through 160,000,000,000 bits per second.

The actual delay buffer is based on the calculations described in tables in the *Class of Service Feature Guide for Routing Devices*. For an example showing how the delay-buffer rates are applied, see [“Example: Configuring Guaranteed Minimum Rate” on page 539](#).

If you do not include the **delay-buffer-rate** statement, the delay-buffer calculation is based on the guaranteed rate, the shaping rate if no guaranteed rate is configured, or the scaled shaping rate if the interface is oversubscribed.

If you do not specify a shaping rate or a guaranteed rate, the logical interface receives a minimal delay-buffer rate and minimal bandwidth equal to 4 MTU-sized packets.

You can configure a rate for the delay buffer that is higher than the guaranteed rate. This can be useful when the traffic flow might not require much bandwidth in general, but in some cases can be bursty and therefore needs a large buffer.

Configuring large buffers on relatively low-speed links can cause packet aging. To help prevent this problem, the software requires that the sum of the delay-buffer rates be less than or equal to the port speed. This restriction does not eliminate the possibility of packet aging, so you should be cautious when using the **delay-buffer-rate** statement. Though some amount of extra buffering might be desirable for burst absorption, delay-buffer rates should not far exceed the service rate of the logical interface.

If you configure delay-buffer rates so that the sum exceeds the port speed, the configured delay-buffer rate is not implemented for the last logical interface that you configure. Instead, that logical interface receives a delay-buffer rate of 0, and a warning

message is displayed in the CLI. If bandwidth becomes available (because another logical interface is deleted or deactivated, or the port speed is increased), the configured delay-buffer-rate is reevaluated and implemented if possible.

If the guaranteed rate of a logical interface cannot be implemented, that logical interface receives a delay-buffer rate of 0, even if the configured delay-buffer rate is within the interface speed. If at a later time the guaranteed rate of the logical interface can be met, the configured delay-buffer rate is reevaluated and if the delay-buffer rate is within the remaining bandwidth, it is implemented.

If any logical interface has a configured guaranteed rate, all other logical interfaces on that port that do not have a guaranteed rate configured receive a delay-buffer rate of 0. This is because the absence of a guaranteed rate configuration corresponds to a guaranteed rate of 0 and, consequently, a delay-buffer rate of 0.

3. To assign a scheduler map to the logical interface, include the **scheduler-map** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

```
[edit class-of-service traffic-control-profiles profile-name]  
scheduler-map map-name;
```

For information about configuring schedulers and scheduler maps, see the *Class of Service Feature Guide for Routing Devices*.

4. To enable large buffer sizes to be configured, include the **q-pic-large-buffer** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]  
q-pic-large-buffer;
```

If you do not include this statement, the delay-buffer size is more restricted. For more information, see the *Class of Service Feature Guide for Routing Devices*.

5. To enable scheduling on logical interfaces, include the **per-unit-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name ]  
per-unit-scheduler;
```

When you include this statement, the maximum number of VLANs supported is 767 on a single-port Gigabit Ethernet IQ PIC. On a two-port Gigabit Ethernet IQ PIC, the maximum number is 383.

6. To apply the traffic-scheduling profile to the logical interface, include the **output-traffic-control-profile** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]  
output-traffic-control-profile profile-name;
```

Example: Configuring Guaranteed Minimum Rate

Two logical interface units, 0 and 1, are provisioned with a guaranteed minimum of 750 Kbps and 500 Kbps, respectively. For logical unit 1, the delay buffer is based on the guaranteed rate setting. For logical unit 0, a delay-buffer rate of 500 Kbps is specified.

The actual delay buffers allocated to each logical interface are 2 seconds of 500 Kbps. The 2-second value is based on the following calculation:

$\text{delay-buffer-rate} < [8 \times 64 \text{ Kbps}])$: 2 seconds of delay-buffer-rate

For more information about this calculation, see the *Class of Service Feature Guide for Routing Devices*.

```
chassis {
  fpc 3 {
    pic 0 {
      q-pic-large-buffer;
    }
  }
}
interfaces {
  tl-3/0/1 {
    per-unit-scheduler;
  }
}
class-of-service {
  traffic-control-profiles {
    tc-profile3 {
      guaranteed-rate 750k;
      scheduler-map sched-map3;
      delay-buffer-rate 500k; # 500 Kbps is less than 8 x 64 Kbps
    }
    tc-profile4 {
      guaranteed-rate 500k; # 500 Kbps is less than 8 x 64 Kbps
      scheduler-map sched-map4;
    }
  }
}
interface tl-3/0/1 {
  unit 0 {
    output-traffic-control-profile tc-profile3;
  }
  unit 1 {
    output-traffic-control-profile tc-profile4;
  }
}
```

**Related
Documentation**

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces on page 561](#)
- [Oversubscribing Interface Bandwidth on LSQ Interfaces on page 532](#)
- [Link Services Configuration for Junos Interfaces](#)

PART 9

Configuring Interface Redundancy and Bundling on LSQ Interfaces

- [Overview on page 543](#)
- [Configuring Interface Redundancy with SONET APS and Virtual Interfaces on page 545](#)
- [Enabling Bundling on LSQ Interfaces on page 559](#)

Overview

- [Layer 2 Service Package Capabilities and Interfaces](#) on page 543

Layer 2 Service Package Capabilities and Interfaces

As described in [“Enabling Service Packages”](#) on page 11, you can configure the AS or Multiservices PIC and the internal ASM in the M7i platform to use either the Layer 2 or the Layer 3 service package.

When you enable the Layer 2 service package, the AS or Multiservices PIC supports *link services*. On the AS or Multiservices PIC and the ASM, link services include the following:

- Junos CoS components—[“Configuring CoS Scheduling Queues on Logical LSQ Interfaces”](#) on page 523 describes how the Junos CoS components work on link services IQ (**lsq**) interfaces. For detailed information about Junos CoS components, see the *Class of Service Feature Guide for Routing Devices*.
- Data compression using the compressed Real-Time Transport Protocol (CRTTP) for use in voice over IP (VoIP) transmission.



NOTE: On LSQ interfaces, all multilink traffic for a single bundle is sent to a single processor. If CRTTP is enabled on the bundle, it adds overhead to the CPU. Because T3 network interfaces support only one link per bundle, make sure you configure a fragmentation map for compressed traffic on these interfaces and specify the **no-fragmentation** option. For more information, see [“Configuring Delay-Sensitive Packet Interleaving”](#) on page 624 and [“Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces”](#) on page 527.

- Link fragment interleaving (LFI) on Frame Relay links using FRF.12 end-to-end fragmentation—The standard for FRF.12 is defined in the specification FRF.12, *Frame Relay Fragmentation Implementation Agreement*.
- LFI on Multilink Point-to-Point Protocol (MLPPP) links.
- Multilink Frame Relay (MLFR) end-to-end (FRF.15)—The standard for FRF.15 is defined in the specification FRF.15, *End-to-End Multilink Frame Relay Implementation Agreement*.

- Multilink Frame Relay (MLFR) UNI NNI (FRF.16)—The standard for FRF.16 is defined in the specification FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*.
- MLPPP—The standard for MLPPP is defined in the specification RFC 1990, *The PPP Multilink Protocol (MP)*.
- Multiclass extension to MLPPP—The standard is defined in the specification RFC 2686, *The Multi-Class Extension to Multi-Link PPP*.

For the LSQ interface on the AS or Multiservices PIC, the configuration syntax is almost the same as for Multilink and Link Services PICs. The primary difference is the use of the interface-type descriptor **lsq** instead of **ml** or **ls**. When you enable the Layer 2 service package on the AS or Multiservices PIC, the following interfaces are automatically created:

```
gr-fpc/pic/port
ip-fpc/pic/port
lsq-fpc/pic/port
lsq-fpc/pic/port:0
...
lsq-fpc/pic/port:N
mt-fpc/pic/port
pd-fpc/pic/port
pe-fpc/pic/port
sp-fpc/pic/port
vt-fpc/pic/port
```

Interface types **gr**, **ip**, **mt**, **pd**, **pe**, and **vt** are standard tunnel interfaces that are available on the AS or Multiservices PIC whether you enable the Layer 2 or the Layer 3 service package. These tunnel interfaces function the same way for both service packages, except that the Layer 2 service package does not support some tunnel functions, as shown in Table 5 on page 24. For more information about tunnel interfaces, see *Tunnel Properties*.



NOTE: Interface type **sp** is created because it is needed by the Junos OS. For the Layer 2 service package, the **sp** interface is not configurable, but you should not disable it.

Interface type **lsq-fpc/pic/port** is the physical link services IQ interface (**lsq**). Interface types **lsq-fpc/pic/port:0** through **lsq-fpc/pic/port:N** represent FRF.16 bundles. These interface types are created when you include the **mlfr-uni-nni-bundles** statement at the **[edit chassis fpc slot-number pic pic-number]** hierarchy level. For more information, see “Configuring CoS Scheduling Queues on Logical LSQ Interfaces” on page 523.



NOTE: On DS0, E1, or T1 interfaces in LSQ bundles, you can configure the **bandwidth** statement, but the router does not use the bandwidth value if the interfaces are included in an MLPPP or MLFR bundle. The bandwidth is calculated internally according to the time slots, framing, and byte-encoding of the interface. For more information about these properties, see the *Junos OS Network Interfaces Library for Routing Devices*.

CHAPTER 39

Configuring Interface Redundancy with SONET APS and Virtual Interfaces

- [Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS on page 545](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using SONET APS on page 548](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 548](#)

Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS

Link services IQ (**lsq-**) interfaces that are paired with SONET PICs can use the Automatic Protection Switching (APS) configuration already available on SONET networks to provide failure recovery. SONET APS provides stateless failure recovery, if it is configured on SONET interfaces in separate chassis and each SONET PIC is paired with an AS or Multiservices PIC in the same chassis. If one of the following conditions for APS failure is met, the associated SONET PIC triggers recovery to the backup circuit and its associated AS or Multiservices PIC. The failure conditions are:

- Failure of Link Services IQ PIC
- Failure of FPC that hosts the Link Services IQ PIC
- Failure of Packet Forwarding Engine
- Failure of chassis

The guidelines for configuring SONET APS are described in the *Junos OS Network Interfaces Library for Routing Devices*.

The following sections describe how to configure failover properties:

- [Configuring the Association between LSQ and SONET Interfaces on page 546](#)
- [Configuring SONET APS Interoperability with Cisco Systems FRF.16 on page 547](#)
- [Restrictions on APS Redundancy for LSQ Interfaces on page 547](#)

Configuring the Association between LSQ and SONET Interfaces

To configure the association between AS or Multiservices PICs hosting link services IQ interfaces and the SONET interfaces, include the **lsq-failure-options** statement at the **[edit interfaces]** hierarchy level:

```
lsq-fpc/pic/port {
  lsq-failure-options {
    no-termination-request;
    [ trigger-link-failure interface-name ];
  }
}
```

For example, consider the following network scenario:

- Primary router includes interfaces **oc3-0/2/0** and **lsq-1/1/0**.
- Backup router includes interfaces **oc3-2/2/0** and **lsq-3/2/0**.

Configure SONET APS, with **oc3-0/2/0** as the working circuit and **oc3-2/2/0** as the protect circuit. Include the **trigger-link-failure** statement to extend failure to the LSQ PICs:

```
interfaces lsq-1/1/0 {
  lsq-failure-options {
    trigger-link-failure oc3-0/2/0;
  }
}
```



NOTE: You must configure the **lsq-failure-options** statement on the primary router only. The configuration is not supported on the backup router.

To inhibit the router from sending PPP termination-request messages to the remote host if the Link Services IQ PIC fails, include the **no-termination-request** statement at the **[edit interfaces lsq-fpc/pic/port lsq-failure-options]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port lsq-failure-options]
no-termination-request;
```

This functionality is supported on link PICs as well. To inhibit the router from sending PPP termination-request messages to the remote host if a link PIC fails, include the **no-termination-request** statement at the **[edit interfaces interface-name ppp-options]** hierarchy level.

```
[edit interfaces interface-name ppp-options]
no-termination-request;
```

The **no-termination-request** statement is supported only with MLPPP and SONET APS configurations and works with PPP, PPP over Frame Relay, and MLPPP interfaces only, on the following PICs:

- Channelized OC3 IQ PICs
- Channelized OC12 IQ PICs

- Channelized STM1 IQ PICs
- Channelized STM4 IQ PICs

Configuring SONET APS Interoperability with Cisco Systems FRF.16

Juniper Networks routers configured with APS might not interoperate correctly with Cisco FRF.16. To enable interoperation, include the **cisco-interoperability** statement at the **[edit interfaces lsq-fpc/pic/port mlfr-uni-nni-bundle-options]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port mlfr-uni-nni-bundle-options]
cisco-interoperability send-lip-remove-link-for-link-reject;
```

The **send-lip-remove-link-for-link-reject** option prompts the router to send a Link Integrity Protocol remove link when it receives an add-link rejection message.

Restrictions on APS Redundancy for LSQ Interfaces

The following restrictions apply to LSQ failure recovery:

- It applies only to Link Services IQ PICs installed in M Series routers, except for M320 routers.
- You must configure the **failure-options** statement on physical LSQ interfaces, not on MLFR channelized units.
- The Link Services IQ PICs must be associated with SONET link PICs. The paired PICs can be installed on different routers or in the same router; in other words, both interchassis and intrachassis recovery are supported
- Failure recovery is stateless; as a result, route flapping and loss of link state is expected in interchassis recovery, requiring PPP renegotiation. In intrachassis recovery, no impact on traffic is anticipated with Routing Engine failover, but PIC failover results in PPP renegotiation.
- The switchover is not revertive: when the original hardware is restored to service, traffic does not automatically revert back to it.
- Normal APS switchover and PIC-triggered APS switchover can be distinguished only by checking the system log messages.



NOTE: When an AS PIC experiences persistent back pressure as a result of high traffic volume for 3 seconds, the condition triggers an automatic core dump and reboot of the PIC to help clear the blockage. A system log message at level LOG_ERR is generated. This mechanism applies to both Layer 2 and Layer 3 service packages.

Related Documentation

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using SONET APS on page 548](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 548](#)

- [Configuring Link Services and CoS on Services PICs on page 529](#)
- [Link Services Configuration for Junos Interfaces](#)

Configuring LSQ Interface Redundancy in a Single Router Using SONET APS

Stateless switchover from one Link Services IQ PIC to another within the same router can be configured by using the SONET APS mechanism described in “[Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS](#)” on page 545. Each Link Services IQ PIC must be associated with a specified SONET link PIC within the same router.



NOTE: For complete intrachassis recovery, including recovery from Routing Engine failover, graceful Routing Engine switchover (GRES) must be enabled on the router. For more information, see the *Junos OS Administration Library for Routing Devices*.

Related Documentation

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 548](#)
- [Configuring Link Services and CoS on Services PICs on page 529](#)
- [Link Services Configuration for Junos Interfaces](#)

Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces

You can configure failure recovery on M Series, MX Series, and T Series routers that have multiple AS or Multiservices PICs and DPCs with **lsq-** interfaces by specifying a virtual LSQ redundancy (**rlsq**) interface in which the primary Link Services IQ PIC is active and a secondary PIC is on standby. If the primary PIC fails, the secondary PIC becomes active, and all LSQ processing is transferred to it. To determine which PIC is currently active, issue the **show interfaces redundancy** command.



NOTE: This configuration does not require the use of SONET APS for failover. Network interfaces that do not support SONET can be used, such as T1 or E1 interfaces.

The following sections provide more information:

- [Configuring Redundant Paired LSQ Interfaces on page 549](#)
- [Restrictions on Redundant LSQ Interfaces on page 550](#)
- [Configuring Link State Replication for Redundant Link PICs on page 551](#)
- [Examples: Configuring Redundant LSQ Interfaces for Failure Recovery on page 553](#)

Configuring Redundant Paired LSQ Interfaces

The physical interface type **rlsq** specifies the pairings between primary and secondary **lsq** interfaces to enable redundancy. To configure a backup **lsq** interface, include the **redundancy-options** statement at the **[edit interfaces rlsqnumber]** hierarchy level:

```
[edit interfaces rlsqnumber]
redundancy-options {
  (hot-standby | warm-standby);
  primary lsq-fpc/pic/port;
  secondary lsq-fpc/pic/port;
}
```

For the **rlsq** interface, **number** can be from 0 through 1023. If the primary **lsq** interface fails, traffic processing switches to the secondary interface. The secondary interface remains active even after the primary interface recovers. If the secondary interface fails and the primary interface is active, processing switches to the primary interface.

The **hot-standby** option is used with one-to-one redundancy configurations, in which one working PIC is supported by one backup PIC. It is supported with MLPPP, CRTTP, FRF.15, and FRF.16 configurations for the LSQ interface to achieve an uninterrupted LSQ service. It sets the requirement for the failure detection and recovery time to be less than 5 seconds. The behavior is revertive, but you can manually switch between the primary and secondary PICs by issuing the **request interfaces (revert | switchover) rlsqnumber** operational mode command. It also provides a switch over time of 5 seconds and less for FRF.15 and a maximum of 10 seconds for FRF.16.

The **warm-standby** option is used with redundancy configurations in which one backup PIC supports multiple working PICs. Recovery times are not guaranteed, because the configuration must be completely restored on the backup PIC after a failure is detected.

Certain combinations of **hot-standby** and **warm-standby** configuration are not permitted and result in a configuration error. The following examples are permitted:

- Interface **rlsq0** configured with **primary lsq-0/0/0** and **warm-standby**, in combination with interface **rlsq0:0** configured with **primary lsq-0/0/0:0**
- Interface **rlsq0:0** configured with **primary lsq-0/0/0:0**, in combination with interface **rlsq0:1** configured with **primary lsq-0/0/0:1**

The following example combinations are not permitted:

- Interface **rlsq0** configured with **primary lsq-0/0/0** and **hot-standby**, in combination with interface **rlsq0:0** configured with **primary lsq-0/0/0:0**
- Interface **rlsq0:0** configured with **primary lsq-0/0/0:0**, in combination with interface **rlsq1:0** configured with **primary lsq-0/0/0:0**
- Interface **rlsq0:0** configured with **primary lsq-0/0/0:1**, in combination with interface **rlsq1:1** configured with **primary lsq-0/0/0:1**
- Interface **rlsq0** configured with **primary lsq-0/0/0**, in combination with interface **rlsq1** configured with **primary lsq-0/0/0**

In addition, the same physical interface cannot be reused as the primary interface for more than one **rlsq** interface, nor can any of the associated logical interfaces. For example, primary interface **lsq-0/0/0** cannot be reused in another **rlsq** interface as **lsq-0/0/0:0**.

Restrictions on Redundant LSQ Interfaces

Link Services IQ PIC failure occurs under the following conditions:

- The primary PIC fails to boot. In this case, the **rlsq** interface does not come up and manual intervention is necessary to reboot or replace the PIC, or to rename the primary PIC to the secondary one in the **rlsq** configuration.
- When configuring an **rlsq** interface, ensure that:
 - The unit number allocated to the **rlsq** interface is less than the number of Multilink Frame Relay user-to-network interface network-to-network interface (UNI-NNI) (FRF.16) bundles allocated on the Link Services PIC.
 - Data-link connection identifier (DLCI) is configured for the **rlsq** interface.

If these conditions are not met, the **rlsq** interface does not boot. When you issue the **show interfaces redundancy** command, the state of the **rlsq** interface is indicated as **Waiting for primary MS PIC**.

- The primary PIC becomes active and then fails. The secondary PIC automatically takes over processing.
- A failover to the secondary PIC takes place. The secondary PIC then fails. If the primary PIC has been restored to active state, processing switches to it.
- The FPC that contains the Link Services IQ PIC fails.

The following constraints apply to redundant LSQ configurations:

- We recommend that primary and secondary PICs be configured in two different FPCs (in chassis other than M10i routers).
- You cannot configure a Link Services IQ PIC with explicit bundle configurations and as a constituent of an **rlsq** interface.
- Redundant LSQ configurations provide full GRES support. (You must configure GRES at the **[edit chassis]** hierarchy level; see the *Junos OS Administration Library for Routing Devices*).
- If you configure the **redundancy-options** statement with the **hot-standby** option, the configuration must include one **primary** interface value and one **secondary** interface value.
- Since the same interface name is used for **hot-standby** and **warm-standby**, if you modify the configuration to change this attribute, it is recommended that you first deactivate the interface, commit the new configuration, and then reactivate the interface.
- You cannot make changes to an active **redundancy-options** configuration. You must deactivate the **rlsqnumber** interface configuration, change it, and reactivate it.

- The **rlsqnumber** configuration becomes active only if the primary interface is active. When the configuration is first activated, the primary interface must be active; if not, the **rlsq** interface waits until the primary interface comes up.
- You cannot modify the configuration of **lsq** interfaces after they have been included in an active **rlsq** interface.
- All the operational mode commands that apply to **rsp** interfaces also apply to **rlsq** interfaces. You can issue **show** commands for the **rlsq** interface or the primary and secondary **lsq** interfaces. However, statistics on the link interfaces are not carried over following a Routing Engine switchover.
- The **rlsq** interfaces also support the **lsq-failure-options** configuration, discussed in [“Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS” on page 545](#). If the primary and secondary Link Services IQ PICs fail and the **lsq-failure-options** statement is configured, the configuration triggers a SONET APS switchover.
- Redundant LSQ configurations that require MLPPP Multilink Frame Relay (FRF.15 and FRF.16) are supported only with the **warm-standby** option.
- Redundant LSQ support is extended to ATM network interfaces.
- Channelized interfaces are used with FRF-16 bundles, for example **rlsq0:0**. The **rlsq** number and its constituents, the **primary** and **secondary** interfaces, must match for the configuration to be valid: either all must be channelized, or none. For an example of an FRF.16 configuration, see [“Configuring LSQ Interface Redundancy for an FRF.16 Bundle” on page 556](#).



NOTE: Adaptive Services and Multiservices PICs in layer-2 mode (running Layer 2 services) are not rebooted when a MAC flow-control situation is detected.

Configuring Link State Replication for Redundant Link PICs

Link state replication, also called *interface preservation*, is an addition to the SONET Automatic Protection Switching (APS) functionality that helps promote redundancy of the link PICs used in LSQ configurations.

Link state replication provides the ability to add two sets of links, one from the active (working) SONET PIC and the other from the backup (protect) SONET PIC to the same bundle. If the active SONET PIC fails, links from the standby PIC are used without causing a link renegotiation. All the negotiated state is replicated from the active links to the standby links to prevent link renegotiation. For more information about SONET APS configurations, see the *Junos OS Network Interfaces Library for Routing Devices*.

To configure link state replication, include the **preserve-interface** statement at the **[edit interfaces interface-name sonet-options aps]** hierarchy level on both network interfaces:

```
edit interfaces interface-name sonet-options aps]
  preserve-interface;
```

The following constraints apply to link PIC redundancy:

- APS functionality must be available on the SONET PICs and the interface configurations must be identical on both ends of the link. Any configuration mismatch causes the commit operation to fail.
- This feature is supported only with LSQ and SONET APS-enabled link PICs, including Channelized OC3, Channelized OC12, and Channelized STM1 intelligent queuing (IQ) PICs.
- Link state replication supports MLPPP and PPP over Frame Relay (**frame-relay-ppp**) encapsulation, and fully supports GRES.
- Enabling the interface or protocol traceoptions with a large number of MLPPP links can trigger Link Control Protocol (LCP) renegotiation during the link switchover time.



NOTE: This renegotiation is more likely to take place for configurations with back-to-back Juniper Networks routers than in networks in which a Juniper Networks router is connected to an add/drop multiplexer (ADM).

- In general, networks that connect a Juniper Networks router to an ADM allow faster MLPPP link switchover than those with back-to-back Juniper Networks routers. The MLPPP link switchover time difference may be significant, especially for networks with a large number of MLPPP links.
- An aggressive LCP keepalive timeout configuration can lead to LCP renegotiation during the MLPPP link switchover. By default, the LCP keepalive timer interval is 10 seconds and the consecutive link down count is 3. The MLPPP links start LCP negotiation only after a timeout of 30 seconds. Lowering these configuration values may trigger one or more of the MLPPP links to renegotiate during the switchover time.



NOTE: LCP renegotiation is more likely to take place for configurations with back-to-back Juniper Networks routers than in networks in which a Juniper Networks router is connected to an ADM.

As an example, the following configuration shows the link state replication configuration between the ports **coc3-1/0/0** and **coc3-2/0/0**.

```
interfaces {
  coc3-1/0/0 {
    sonet-options {
      aps {
        preserve-interface;
        working-circuit aps-group-1;
      }
    }
  }
  coc3-2/0/0 {
    sonet-options {
      aps {
        preserve-interface;
      }
    }
  }
}
```

```

        protect-circuit aps-group-1;
    }
}
}

```

Examples: Configuring Redundant LSQ Interfaces for Failure Recovery

Configuring LSQ Interface Redundancy for MLPPP

The following configuration shows that **lsq-1/1/0** and **lsq-1/3/0** work as a pair and the redundancy type is **hot-standby**, which sets the requirement for the failure detection and recovery time to be less than 5 seconds:

```

interfaces rlsq0 {
  redundancy-options {
    primary lsq-1/1/0;
    secondary lsq-1/3/0;
    hot-standby; #either hot-standby or warm-standby is supported
  }
}

```

The following example shows a related MLPPP configuration:



NOTE: MLPPP protocol configuration is required for this configuration.

```

interfaces {
  t1-1/1/2/0 {
    unit 0 {
      family mlppp {
        bundle rlsq0.0;
      }
    }
  }
  rlsq0 {
    unit 0 {
      family inet {
        address 30.1.1.2/24;
      }
    }
  }
}

```

The following example shows a related CoS configuration:

```

class-of-service {
  interfaces {
    rlsq0 {
      unit * {
        fragmentation-maps fr-map1;
      }
    }
  }
}

```

```
}

```

The following example shows a complete link state replication configuration for MLPPP. This example uses two bundles, each with four T1 links. The first four T1 links (**t1-*:1** through **t1-*:4**) form the first bundle and the last four T1 links (**t1-*:5** through **t1-*:8**) form the second bundle. To minimize the duplication in the configuration, this example uses the **[edit groups]** statement; for more information, see the *Junos OS Administration Library for Routing Devices*. This type of configuration is not required; it simplifies the task and minimizes duplication.

```
groups {
  ml-partition-group {
    interfaces {
      <coc3-*> {
        partition 1 oc-slice 1 interface-type coc1;
      }
      <coc1-*> {
        partition 1-8 interface-type t1;
      }
    }
  }
  ml-bundle-group-1 {
    interfaces {
      <t1-*:"[1-4]"> {
        encapsulation ppp;
        unit 0 {
          family mlppp {
            bundle lsq-0/1/0.0;
          }
        }
      }
    }
  }
  ml-bundle-group-2 {
    interfaces {
      <t1-*:"[5-8]"> {
        encapsulation ppp;
        unit 0 {
          family mlppp {
            bundle lsq-0/1/0.1;
          }
        }
      }
    }
  }
}
interfaces {
  lsq-0/1/0 {
    unit 0 {
      encapsulation multilink-ppp;
      family inet {
        address 1.1.1/32 {
          destination 1.1.1.2;
        }
      }
    }
  }
}
```

```
}
unit 1 {
    encapsulation multilink-ppp;
    family inet {
        address 1.1.2.1/32 {
            destination 1.1.2.2;
        }
    }
}
}
coc3-1/0/0 {
    apply-groups ml-partition-group;
    sonet-options {
        aps {
            preserve-interface;
            working-circuit aps-group-1;
        }
    }
}
coc1-1/0/0:1 {
    apply-groups ml-partition-group;
}
t1-1/0/0:1:1 {
    apply-groups ml-bundle-group-1;
}
t1-1/0/0:1:2 {
    apply-groups ml-bundle-group-1;
}
t1-1/0/0:1:3 {
    apply-groups ml-bundle-group-1;
}
t1-1/0/0:1:4 {
    apply-groups ml-bundle-group-1;
}
t1-1/0/0:1:5 {
    apply-groups ml-bundle-group-2;
}
t1-1/0/0:1:6 {
    apply-groups ml-bundle-group-2;
}
t1-1/0/0:1:7 {
    apply-groups ml-bundle-group-2;
}
t1-1/0/0:1:8 {
    apply-groups ml-bundle-group-2;
}
coc3-2/0/0 {
    apply-groups ml-partition-group;
    sonet-options {
        aps {
            preserve-interface;
            protect-circuit aps-group-1;
        }
    }
}
coc1-2/0/0:1 {
```

```
        apply-groups ml-partition-group;
    }
    t1-2/0/0:1:1 {
        apply-groups ml-bundle-group-1;
    }
    t1-2/0/0:1:2 {
        apply-groups ml-bundle-group-1;
    }
    t1-2/0/0:1:3 {
        apply-groups ml-bundle-group-1;
    }
    t1-2/0/0:1:4 {
        apply-groups ml-bundle-group-1;
    }
    t1-2/0/0:1:5 {
        apply-groups ml-bundle-group-2;
    }
    t1-2/0/0:1:6 {
        apply-groups ml-bundle-group-2;
    }
    t1-2/0/0:1:7 {
        apply-groups ml-bundle-group-2;
    }
    t1-2/0/0:1:8 {
        apply-groups ml-bundle-group-2;
    }
}
```

Configuring LSQ Interface Redundancy for an FRF.15 Bundle

The following example shows a configuration for an FRF.15 bundle:

```
interfaces rlsq0 {
    redundancy-options {
        primary lsq-1/2/0;
        secondary lsq-1/3/0;
        warm-standby; #either hot-standby or warm-standby is supported
    }
    unit 0 {
        encapsulation multilink-frame-relay-end-to-end;
        family inet {
            address 30.1.1.1/24;
        }
    }
}
```

Configuring LSQ Interface Redundancy for an FRF.16 Bundle

The following example shows a configuration for an FRF.16 bundle:

```
interfaces rlsq0:0 {
    dce;
    encapsulation multilink-frame-relay-uni-nni;
```

```
redundancy-options {  
  primary lsq-1/2/0:0;  
  secondary lsq-1/3/0:0;  
  warm-standby; #either hot-standby or warm-standby is supported  
}  
unit 0 {  
  dlc 1000;  
  family inet {  
    address 50.1.1.1/24;  
  }  
}
```

**Related
Documentation**

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS on page 545](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using SONET APS on page 548](#)
- [Configuring Link Services and CoS on Services PICs on page 529](#)
- *Link Services Configuration for Junos Interfaces*

CHAPTER 40

Enabling Bundling on LSQ Interfaces

- [Inline MLPPP for WAN Interfaces Overview on page 559](#)
- [Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces on page 561](#)
- [Configuring Multiclass MLPPP on LSQ Interfaces on page 562](#)
- [Enabling Inline LSQ Services on page 563](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using MLPPP on page 565](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.16 on page 571](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.15 on page 576](#)
- [Configuring LSQ Interfaces for Single Fractional T1 or El Interfaces Using MLPPP and LFI on page 577](#)
- [Configuring LSQ Interfaces for Single Fractional T1 or El Interfaces Using FRF.12 on page 582](#)
- [Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 589](#)
- [Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12 on page 591](#)
- [Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP on page 593](#)

Inline MLPPP for WAN Interfaces Overview

Inline Multilink PPP (MLPPP), Multilink Frame Relay (FRF.16), and Multilink Frame Relay End-to-End (FRF.15) for time-division multiplexing (TDM) WAN interfaces provide bundling services through the Packet Forwarding Engine without requiring a PIC or Dense Port Concentrator (DPC).

Traditionally, bundling services are used to bundle multiple low-speed links to create a higher bandwidth pipe. This combined bandwidth is available to traffic from all links and supports link fragmentation and interleaving (LFI) on the bundle, reducing high priority packet transmission delay.

This support includes multiple links on the same bundle as well as multiclass extension for MLPPP. Through this service you can enable bundling services without additional DPC slots to support Service DPC and free up the slots for other MICs.



NOTE: MLPPP is not supported on MX Series Virtual Chassis.

Configuring inline MLPPP for WAN interfaces benefits the following services:

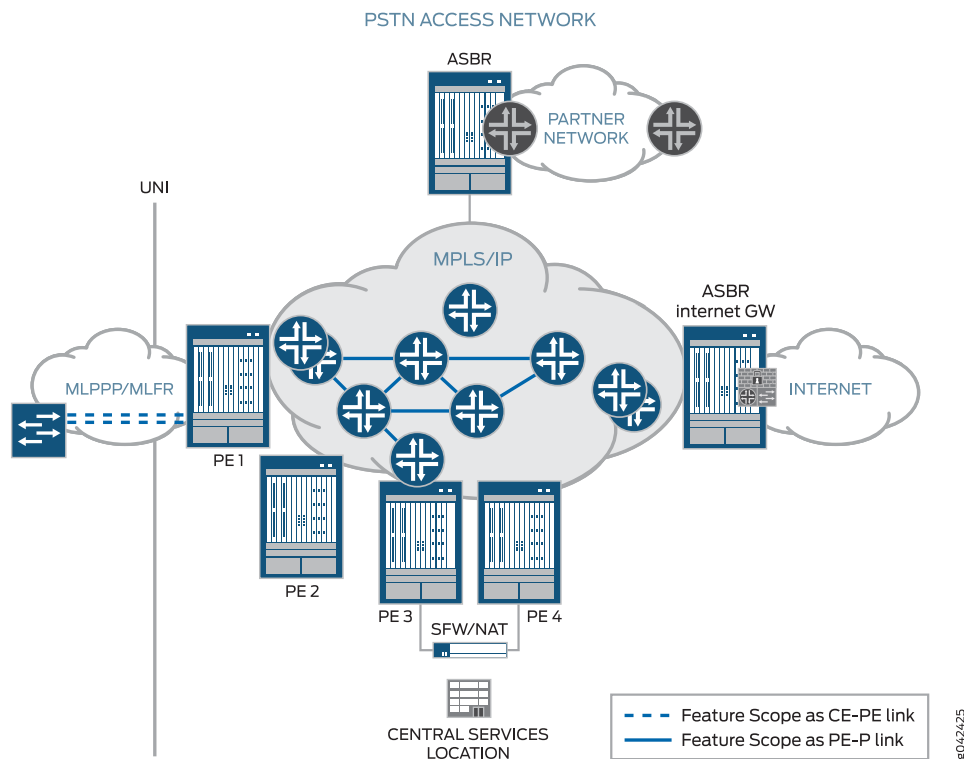
- CE-PE link for Layer 3 VPN and DIA service with public switched telephone networks (PSTN)-based access networks.
- PE-P link when PSTN is used for MPLS networks.

This feature is used by the following service providers:

- Service providers that use PSTN to offer Layer 3 VPN and DIA service with PSTN-based access networks to medium or large business customers.
- Service providers with SONET-based core networks.

The following figure illustrates the scope of this feature:

Figure 25: Inline MLPPP for WAN Interfaces



For connecting many smaller sites in VPNs, bundling the TDM circuits together with MLPPP/MLFR technology is the *only* way to offer higher bandwidth and link redundancy.

MLPPP enables you to bundle multiple PPP links into a single multilink bundle, and MLFR enables you to bundle multiple Frame Relay data-link connection identifiers (DLCIs) into

a single multilink bundle. Multilink bundles provide additional bandwidth, load balancing, and redundancy by aggregating low-speed links, such as T1, E1, and serial links.

MLPPP is a protocol for aggregating multiple constituent links into one larger PPP bundle. MLFR allows you to aggregate multiple Frame Relay links by inverse multiplexing. MLPPP and MLFR provide service options between low-speed T1 and E1 services. In addition to providing additional bandwidth, bundling multiple links can add a level of fault tolerance to your dedicated access service. Because you can implement bundling across multiple interfaces, you can protect users against loss of access when a single interface fails.

To configure inline MLPPP for WAN interfaces, see:

- [Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End \(FRF.15\) for WAN Interfaces on page 754](#)
- [Example: Configuring Inline Multilink Frame Relay \(FRF.16\) for WAN Interfaces on page 788](#)

Related Documentation

- [Configuring the Junos OS to Support the Link Services PIC](#)
- [Enabling Inline LSQ Services on page 563](#)
- [Enabling MLPPP Link Fragmentation and Interleaving on page 769](#)
- [Example: Configuring Multilink Frame Relay FRF.15 on page 783](#)
- [Example: Configuring Multilink Frame Relay FRF.16 on page 779](#)
- [Link and Multilink Services Interfaces Feature Guide for Routing Devices](#)

Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces

Link-layer overhead can cause packet drops on constituent links because of bit stuffing on serial links. Bit stuffing is used to prevent data from being interpreted as control information.

By default, 4 percent of the total bundle bandwidth is set aside for link-layer overhead. In most network environments, the average link-layer overhead is 1.6 percent. Therefore, we recommend 4 percent as a safeguard. For more information, see RFC 4814, *Hash and Stuffing: Overlooked Factors in Network Device Benchmarking*.

For link services IQ (**lsq-**) interfaces, you can configure the percentage of bundle bandwidth to be set aside for link-layer overhead. To do this, include the **link-layer-overhead** statement:

```
link-layer-overhead percent;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* *mlfr-uni-nni-bundle-options*]**
- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

You can configure the value to be from 0 percent through 50 percent.

**Related
Documentation**

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Oversubscribing Interface Bandwidth on LSQ Interfaces on page 532](#)
- [Configuring Guaranteed Minimum Rate on LSQ Interfaces on page 537](#)
- [Link Services Configuration for Junos Interfaces](#)

Configuring Multiclass MLPPP on LSQ Interfaces

For link services IQ (**lsq**-) interfaces with MLPPP encapsulation, you can configure multiclass MLPPP (MCML). If you do not configure MCML, fragments from different classes cannot be interleaved. All fragments for a single packet must be sent before the fragments from another packet are sent. Nonfragmented packets can be interleaved between fragments of another packet to reduce latency seen by nonfragmented packets. In effect, latency-sensitive traffic is encapsulated as regular PPP traffic, and bulk traffic is encapsulated as multilink traffic. This model works as long as there is a single class of latency-sensitive traffic, and there is no high-priority traffic that takes precedence over latency-sensitive traffic. This approach to LFI, used on the Link Services PIC, supports only two levels of traffic priority, which is not sufficient to carry the four-to-eight forwarding classes that are supported by M Series and T Series routers. For more information about the Link Services PIC support of LFI, see “[Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces](#)” on page 773.

For link services IQ interfaces only, you can configure MCML, as defined in RFC 2686, *The Multi-Class Extension to Multi-Link PPP*. MCML makes it possible to have multiple classes of latency-sensitive traffic that are carried over a single multilink bundle with bulk traffic. In effect, MCML allows different classes of traffic to have different latency guarantees. With MCML, you can map each forwarding class into a separate multilink class, thus preserving priority and latency guarantees.



NOTE: Configuring both LFI and MCML on the same bundle is not necessary, nor is it supported, because multiclass MLPPP represents a superset of functionality. When you configure multiclass MLPPP, LFI is automatically enabled.

The Junos OS implementation of MCML does not support compression of common header bytes, which is referred to in RFC 2686 as “prefix elision.”

MCML greatly simplifies packet ordering issues that occur when multiple links are used. Without MCML, all voice traffic belonging to a single flow is hashed to a single link to avoid packet ordering issues. With MCML, you can assign voice traffic to a high-priority class, and you can use multiple links. For more information about voice services support on link services IQ interfaces (**lsq**), see “[Configuring Services Interfaces for Voice Services](#)” on page 622.

To configure MCML on a link services IQ interface, you must specify how many multilink classes should be negotiated when a link joins the bundle, and you must specify the mapping of a forwarding class into an MCML class.

To specify how many multilink classes should be negotiated when a link joins the bundle, include the **multilink-max-classes** statement:

```
multilink-max-classes number;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

The number of multilink classes can be 1 through 8. The number of multilink classes for each forwarding class must not exceed the number of multilink classes to be negotiated.

To specify the mapping of a forwarding class into a MCML class, include the **multilink-class** statement at the **[edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]** hierarchy level:

```
[edit class-of-service fragmentation-maps map-name forwarding-class class-name]  
multilink-class number;
```

The multilink class index number can be 0 through 7. The **multilink-class** statement and **no-fragmentation** statements are mutually exclusive.

To view the number of multilink classes negotiated, issue the **show interfaces *lsq-fpc/pic/port.logical-unit-number* detail** command.

Related Documentation

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using MLPPP on page 565](#)
- [Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP on page 593](#)
- [Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 589](#)
- [Link Services Configuration for Junos Interfaces](#)

Enabling Inline LSQ Services

Inline Multilink PPP (MLPPP), Multilink Frame Relay (FRF.16), and Multilink Frame Relay End-to-End (FRF.15) for time-division multiplexing (TDM) WAN interfaces provide bundling services through the Packet Forwarding Engine without requiring a PIC or Dense Port Concentrator (DPC).

Traditionally, bundling services are used to bundle multiple low-speed links to create a higher bandwidth pipe. This combined bandwidth is available to traffic from all links and supports link fragmentation and interleaving (LFI) on the bundle, reducing high priority packet transmission delay.

This support includes multiple links on the same bundle as well as multiclass extension for MLPPP. Through this service you can enable bundling services without additional DPC slots to support Service DPC and free up the slots for other MICs.

The inline LSQ logical interface (referred to as `lsq-`) is a virtual service logical interface that resides on the Packet Forwarding Engine to provide Layer 2 bundling services that do not need a service PIC. The naming convention is `lsq-slot/pic/O`. Currently, only TYPE1 and TYPE2 queuing Modular Point Concentrators (MPCs) support inline LSQ logical interfaces. A Type1 MPC has only one logical unit (LU); therefore only one LSQ logical interface can be created. When configuring a Type1 MPC, use PIC slot 0. Type2 MPC has two LUs; therefore two LSQ logical interfaces can be created. When configuring a Type2 MPC, use PIC slot 0 and slot 2.

Configure each LSQ logical interface with one loopback stream. This stream can be shaped like a regular stream, and is shared with other inline interfaces, such as the inline services (SI) interface.

To support FRF.16 bundles, create logical interfaces with the naming convention `lsq-slot/pic/O:bundle_id`, where *bundle_id* can range from 0 to 254. You can configure logical interfaces created on the main LSQ logical interface as MLPPP or FRF.16.

Because SI and LSQ logical interfaces might share the same stream, and there could be multiple LSQ logical interfaces on that stream, any logical interface-related shaping is configured at the Layer 2 node instead of the Layer 1 node. As a result, when SI is enabled, instead of limiting the stream bandwidth to 1Gb or 10Gb based on the configuration, only the Layer 2 queue allocated for the SI interface is shaped at 1Gb or 10Gb.

For MLPPP and FRF.15, each LSQ logical interface is shaped based on the total bundle bandwidth (sum of member link bandwidths with control packet flow overhead) by configuring one unique Layer 3 node per bundle. Similarly, each FRF.16 logical interface is shaped based on total bundle bandwidth by configuring one unique Layer 2 node per bundle. FRF.16 logical interface data-link connection identifiers (DLCIs) are mapped to Layer 3 nodes.

To enable inline LSQ services and create the `lsq-` logical interface for the specified PIC, specify the `multi-link-layer-2-inline` and `mlfr-uni-nni-bundles-inline` configuration statements.

```
[edit chassis fpc number pic number]
user@host# set multi-link-layer-2-inline
user@host# set mlfr-uni-nni-bundles-inline number
```

For example, to enable inline service for PIC 0 on a Type1 MPC on slot 1:

```
[edit chassis fpc 1 pic 0]
user@host# set multi-link-layer-2-inline
user@host# set mlfr-uni-nni-bundles-inline 1
```

As a result, logical interfaces `lsq-1/0/0`, `lsq-1/0/0:0`, and `lsq-1/0/0:1` are created. The number of inline multilink frame relay user-to-network interface (UNI) and network-to-network interface (NNI) bundles is set to 1.

For example, to enable inline service for both PIC 0 and PIC 2 on Type2 MPC installed in slot 5:

```
[edit chassis fpc 5 pic 0]
user@host# set multi-link-layer-2-inline
user@host# set mlfr-uni-nni-bundles-inline 1
```

```
[edit chassis fpc 5 pic 2]
user@host# set multi-link-layer-2-inline
user@host# set mlfr-uni-nni-bundles-inline 1
```

As a result, logical interfaces lsq-5/0/0, lsq-5/0/0:0, lsq-5/0/0:1, lsq-5/2/0, lsq-5/2/0:0, and lsq-5/2/0:1 are created. The number of inline multilink frame relay user-to-network interface (UNI) and network-to-network interface (NNI) bundles is set to 1.



NOTE: The PIC number here is only used as an anchor to choose the correct LU to bind the inline LSQ interface. The bundling services are operational as long as the Packet Forwarding Engine to which it is bound is operational, even if the logical PIC is offline.

Related Documentation

- [Inline MLPPP for WAN Interfaces Overview on page 559](#)
- [Link Services IQ Interfaces](#)
- [Link and Multilink Services Interfaces Feature Guide for Routing Devices](#)
- [mlfr-uni-nni-bundles-inline on page 1412](#)
- [multi-link-layer-2-inline on page 1414](#)

Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP

To configure an NxT1 bundle using MLPPP, you aggregate *N* different T1 links into a bundle. The NxT1 bundle is called a logical interface, because it can represent, for example, a routing adjacency. To aggregate T1 links into a an MLPPP bundle, include the **bundle** statement at the `[edit interfaces t1-fpc/pic/port unit logical-unit-number family mlppp]` hierarchy level:

```
[edit interfaces t1-fpc/pic/port unit logical-unit-number family mlppp]
bundle lsq-fpc/pic/port.logical-unit-number;
```



NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure the link services IQ interface properties, include the following statements at the `[edit interfaces lsq-fpc/pic/port unit logical-unit-number]` hierarchy level:

```
[edit interfaces lsq-fpc/pic/port unit logical-unit-number]
drop-timeout milliseconds;
```

```

encapsulation multilink-ppp;
fragment-threshold bytes;
link-layer-overhead percent;
minimum-links number;
mrru bytes;
short-sequence;
family inet {
    address address;
}

```

The logical link services IQ interface represents the MLPPP bundle. For the MLPPP bundle, there are four associated queues on M Series routers and eight associated queues on M320 and T Series routers. A scheduler removes packets from the queues according to a scheduling policy. Typically, you designate one queue to have strict priority, and the remaining queues are serviced in proportion to weights you configure.

For MLPPP, assign a single scheduler map to the link services IQ interface (**lsq**) and to each constituent link. The default schedulers for M Series and T Series routers, which assign 95, 0, 0, and 5 percent bandwidth for the transmission rate and buffer size of queues 0, 1, 2, and 3, are not adequate when you configure LFI or multiclass traffic. Therefore, for MLPPP, you should configure a single scheduler with nonzero percent transmission rates and buffer sizes for queues 0 through 3, and assign this scheduler to the link services IQ interface (**lsq**) and to each constituent link, as shown in [“Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP” on page 568](#).



NOTE: For M320 and T Series routers, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

If the member link belonging to one MLPP, MLFR, or MFR bundle interface is moved to another bundle interface, or the links are swapped between two bundle interfaces, a commit is required between the delete and add operations to ensure that the configuration is applied correctly.

If the bundle has more than one link, you must include the **per-unit-scheduler** statement at the **[edit interfaces lsq-fpc/pic/port]** hierarchy level:

```

[edit interfaces lsq-fpc/pic/port]
per-unit-scheduler;

```

To configure and apply the scheduling policy, include the following statements at the **[edit class-of-service]** hierarchy level:

```

[edit class-of-service]
interfaces {
    t1-fpc/pic/port unit logical-unit-number {
        scheduler-map map-name;
    }
}
forwarding-classes {
    queue queue-number class-name;
}

```



```

scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
schedulers {
  scheduler-name {
    buffer-size (percent percentage | remainder | temporal microseconds);
    priority priority-level;
    transmit-rate (rate | percent percentage | remainder) <exact>;
  }
}

```

For link services IQ interfaces, a strict-high-priority queue might starve the other three queues because traffic in a strict-high priority queue is transmitted before any other queue is serviced. This implementation is unlike the standard Junos CoS implementation in which a strict-high-priority queue does round-robin with high-priority queues, as described in the *Class of Service Feature Guide for Routing Devices*.

After the scheduler removes a packet from a queue, a certain action is taken. The action depends on whether the packet came from a multilink encapsulated queue (fragmented and sequenced) or a nonencapsulated queue (hashed with no fragmentation). Each queue can be designated as either multilink encapsulated or nonencapsulated, independently of the other. By default, traffic in all forwarding classes is multilink encapsulated. To configure packet fragmentation handling on a queue, include the **fragmentation-maps** statement at the **[edit class-of-service]** hierarchy level:

```

fragmentation-maps {
  map-name {
    forwarding-class class-name {
      fragment-threshold bytes;
      multilink-class number;
      no-fragmentation;
    }
  }
}

```

For NxT1 bundles using MLPPP, the byte-wise load balancing used in multilink-encapsulated queues is superior to the flow-wise load balancing used in nonencapsulated queues. All other considerations are equal. Therefore, we recommend that you configure all queues to be multilink encapsulated. You do this by including the **fragment-threshold** statement in the configuration. If you choose to set traffic on a queue to be nonencapsulated rather than multilink encapsulated, include the **no-fragmentation** statement in the fragmentation map. You use the **multilink-class** statement to map a forwarding class into a multiclass MLPPP (MCML). For more information about MCML, see [“Configuring Multiclass MLPPP on LSQ Interfaces” on page 562](#). For more information about fragmentation maps, see [“Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces” on page 527](#).

When a packet is removed from a multilink-encapsulated queue, the software gives the packet an MLPPP header. The MLPPP header contains a sequence number field, which is filled with the next available sequence number from a counter. The software then

places the packet on one of the N different T1 links. The link is chosen on a packet-by-packet basis to balance the load across the various T1 links.

If the packet exceeds the minimum link MTU, or if a queue has a fragment threshold configured at the **[edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]** hierarchy level, the software splits the packet into two or more fragments, which are assigned consecutive multilink sequence numbers. The outgoing link for each fragment is selected independently of all other fragments.

If you do not include the **fragment-threshold** statement in the fragmentation map, the fragmentation threshold you set at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level is the default for all forwarding classes. If you do not set a maximum fragment size anywhere in the configuration, packets are fragmented if they exceed the smallest MTU of all the links in the bundle.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the maximum received reconstructed unit (MRRU) by including the **mrru** statement at the **[edit interfaces *lsq-fpc/pic/port* unit *logical-unit-number*]** hierarchy level. The MRRU is similar to the MTU, but is specific to link services interfaces. By default the MRRU size is 1500 bytes, and you can configure it to be from 1500 through 4500 bytes. For more information, see [“Configuring MRRU on Multilink and Link Services Logical Interfaces” on page 734](#).

When a packet is removed from a nonencapsulated queue, it is transmitted with a plain PPP header. Because there is no MLPPP header, there is no sequence number information. Therefore, the software must take special measures to avoid packet reordering. To avoid packet reordering, the software places the packet on one of the N different T1 links. The link is determined by hashing the values in the header. For IP, the software computes the hash based on source address, destination address, and IP protocol. For MPLS, the software computes the hash based on up to five MPLS labels, or four MPLS labels and the IP header.

For UDP and TCP the software computes the hash based on the source and destination ports, as well as source and destination IP addresses. This guarantees that all packets belonging to the same TCP/UDP flow always pass through the same T1 link, and therefore cannot be reordered. However, it does not guarantee that the load on the various T1 links is balanced. If there are many flows, the load is usually balanced.

The N different T1 interfaces link to another router, which can be from Juniper Networks or another vendor. The router at the far end gathers packets from all the T1 links. If a packet has an MLPPP header, the sequence number field is used to put the packet back into sequence number order. If the packet has a plain PPP header, the software accepts the packet in the order in which it arrives and makes no attempt to reassemble or reorder the packet.

Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP

```
[edit chassis]
fpc 1 {
  pic 3 {
    adaptive-services {
      service-package layer-2;
```

```

    }
  }
}
[edit interfaces]
t1-0/0/0 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle lsq-1/3/0.1; # This adds t1-0/0/0 to the specified bundle.
    }
  }
}
t1-0/0/1 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle lsq-1/3/0.1;
    }
  }
}
lsq-1/3/0 {
  unit 1 { # This is the virtual link that concatenates multiple T1s.
    encapsulation multilink-ppp;
    drop-timeout 1000;
    fragment-threshold 128;
    link-layer-overhead 0.5;
    minimum-links 2;
    mrru 4500;
    short-sequence;
    family inet {
      address 10.2.3.4/24;
    }
  }
}
[edit interfaces]
lsq-1/3/0 {
  per-unit-scheduler;
}
[edit class-of-service]
interfaces {
  lsq-1/3/0 { # multilink PPP constituent link
    unit 0 {
      scheduler-map sched-map1;
    }
  }
  t1-0/0/0 { # multilink PPP constituent link
    unit 0 {
      scheduler-map sched-map1;
    }
  }
  t1-0/0/1 { # multilink PPP constituent link
    unit 0 {
      scheduler-map sched-map1;
    }
  }
forwarding-classes {
  queue 0 be;
  queue 1 ef;
  queue 2 af;
}

```

```
    queue 3 nc;
}
scheduler-maps {
  sched-map1 {
    forwarding-class af scheduler af-scheduler;
    forwarding-class be scheduler be-scheduler;
    forwarding-class ef scheduler ef-scheduler;
    forwarding-class nc scheduler nc-scheduler;
  }
}
schedulers {
  af-scheduler {
    transmit-rate percent 30;
    buffer-size percent 30;
    priority low;
  }
  be-scheduler {
    transmit-rate percent 25;
    buffer-size percent 25;
    priority low;
  }
  ef-scheduler {
    transmit-rate percent 40;
    buffer-size percent 40;
    priority strict-high; # voice queue
  }
  nc-scheduler {
    transmit-rate percent 5;
    buffer-size percent 5;
    priority high;
  }
}
fragmentation-maps {
  fragmap-1 {
    forwarding-class be {
      fragment-threshold 180;
    }
    forwarding-class ef {
      fragment-threshold 100;
    }
  }
}
[edit interfaces]
lsq-1/3/0 {
  unit 0 {
    fragmentation-map fragmap-1;
  }
}
```

Related Documentation

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.16 on page 571](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.15 on page 576](#)
- [Link Services Configuration for Junos Interfaces](#)

Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16

To configure an NxT1 bundle using FRF.16, you aggregate *N* different T1 links into a bundle. The NxT1 bundle carries a potentially large number of Frame Relay PVCs, identified by their DLCIs. Each DLCI is called a logical interface, because it can represent, for example, a routing adjacency.

To aggregate T1 links into an FRF.16 bundle, include the **mlfr-uni-nni-bundles** statement at the **[edit chassis fpc slot-number pic slot-number]** hierarchy level and include the **bundle** statement at the **[edit interfaces t1-fpc/pic/port unit logical-unit-number family mlfr-uni-nni]** hierarchy level:

```
[edit chassis fpc slot-number pic slot-number]
mlfr-uni-nni-bundles number;
```

```
[edit interfaces t1-fpc/pic/port unit logical-unit-number family mlfr-uni-nni]
bundle lsq-fpc/pic/port:channel;
```



NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure the link services IQ interface properties, include the following statements at the **[edit interfaces lsq- fpc/pic/port:channel]** hierarchy level:

```
[edit interfaces lsq- fpc/pic/port:channel]
encapsulation multilink-frame-relay-uni-nni;
dce;
mlfr-uni-nni-options {
  acknowledge-retries number;
  acknowledge-timer milliseconds;
  action-red-differential-delay (disable-tx | remove-link);
  drop-timeout milliseconds;
  fragment-threshold bytes;
  hello-timer milliseconds;
  link-layer-overhead percent;
  lmi-type (ansi | itu);
  minimum-links number;
  mrru bytes;
  n391 number;
  n392 number;
  n393 number;
  red-differential-delay milliseconds;
  t391 number;
  t392 number;
  yellow-differential-delay milliseconds;
}
unit logical-unit-number {
  dlcid dlcid-identifier;
  family inet {
    address address;
```

```
}  
}
```

The link services IQ channel represents the FRF.16 bundle. Four queues are associated with each DLCI. A scheduler removes packets from the queues according to a scheduling policy. On the link services IQ interface, you typically designate one queue to have strict priority. The remaining queues are serviced in proportion to weights you configure.

For link services IQ interfaces, a strict-high-priority queue might starve the other three queues because traffic in a strict-high-priority queue is transmitted before any other queue is serviced. This implementation is unlike the standard Junos CoS implementation in which a strict-high-priority queue does round-robin with high-priority queues, as described in the *Class of Service Feature Guide for Routing Devices*.

If the bundle has more than one link, you must include the **per-unit-scheduler** statement at the **[edit interfaces lsq-fpc/pic/port:channel]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port:channel]  
  per-unit-scheduler;
```

For FRF.16, you can assign a single scheduler map to the link services IQ interface (**lsq**) and to each link services IQ DLCI, or you can assign different scheduler maps to the various DLCIs of the bundle, as shown in [“Example: Configuring an LSQ Interface as an NxT1 Bundle Using FRF.16” on page 574](#).

For the constituent links of an FRF.16 bundle, you do not need to configure a custom scheduler. Because LFI and multiclass are not supported for FRF.16, the traffic from each constituent link is transmitted from queue 0. This means you should allow most of the bandwidth to be used by queue 0. For M Series and T Series routers, the default schedulers' transmission rate and buffer size percentages for queues 0 through 3 are 95, 0, 0, and 5 percent. These default schedulers send all user traffic to queue 0 and all network-control traffic to queue 3, and therefore are well suited to the behavior of FRF.16. If desired, you can configure a custom scheduler that explicitly replicates the 95, 0, 0, and 5 percent queuing behavior, and apply it to the constituent links.



NOTE: For M320 and T Series routers, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

If the member link belonging to one MLPP, MLFR, or MFR bundle interface is moved to another bundle interface, or the links are swapped between two bundle interfaces, a commit is required between the delete and add operations to ensure that the configuration is applied correctly.

To configure and apply the scheduling policy, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]  
  interfaces {  
    lsq-fpc/pic/port:channel {  
      unit logical-unit-number {
```

```

        scheduler-map map-name;
    }
}
forwarding-classes {
    queue queue-number class-name;
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
schedulers {
    scheduler-name {
        buffer-size (percent percentage | remainder | temporal microseconds);
        priority priority-level;
        transmit-rate (rate | percent percentage | remainder) <exact>;
    }
}

```

To configure packet fragmentation handling on a queue, include the **fragmentation-maps** statement at the **[edit class-of-service]** hierarchy level:

```

[edit class-of-service]
fragmentation-maps {
    map-name {
        forwarding-class class-name {
            fragment-threshold bytes;
        }
    }
}

```

For FRF.16 traffic, only multilink encapsulated (fragmented and sequenced) queues are supported. This is the default queuing behavior for all forwarding classes. FRF.16 does not allow for nonencapsulated traffic because the protocol requires that all packets carry the fragmentation header. If a large packet is split into multiple fragments, the fragments must have consecutive sequential numbers. Therefore, you cannot include the **no-fragmentation** statement at the **[edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]** hierarchy level for FRF.16 traffic. For FRF.16, if you want to carry voice or any other latency-sensitive traffic, you should not use slow links. At T1 speeds and above, the serialization delay is small enough so that you do not need to use explicit LFI.

When a packet is removed from a multilink-encapsulated queue, the software gives the packet an FRF.16 header. The FRF.16 header contains a sequence number field, which is filled with the next available sequence number from a counter. The software then places the packet on one of the *N* different T1 links. The link is chosen on a packet-by-packet basis to balance the load across the various T1 links.

If the packet exceeds the minimum link MTU, or if a queue has a fragment threshold configured at the **[edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]** hierarchy level, the software splits the packet into two or more fragments, which are assigned consecutive multilink sequence numbers. The outgoing link for each fragment is selected independently of all other fragments.

If you do not include the **fragment-threshold** statement in the fragmentation map, the fragmentation threshold you set at the **[edit interfaces *interface-name* unit *logical-unit-number*]** or **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]** hierarchy level is the default for all forwarding classes. If you do not set a maximum fragment size anywhere in the configuration, packets are fragmented if they exceed the smallest MTU of all the links in the bundle.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the maximum received reconstructed unit (MRRU) by including the **mrru** statement at the **[edit interfaces *lsq-fpc/port* unit *logical-unit-number*]** or **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]** hierarchy level. The MRRU is similar to the MTU but is specific to link services interfaces. By default, the MRRU size is 1500 bytes, and you can configure it to be from 1500 through 4500 bytes. For more information, see [“Configuring MRRU on Multilink and Link Services Logical Interfaces” on page 734](#).

The *N* different T1 interfaces link to another router, which can be from Juniper Networks or another vendor. The router at the far end gathers packets from all the T1 links. Because each packet has an FRF.16 header, the sequence number field is used to put the packet back into sequence number order.

Example: Configuring an LSQ Interface as an NxT1 Bundle Using FRF.16

Configure an NxT1 bundle using FRF.16 with multiple CoS scheduler maps:

```
[edit chassis fpc 1 pic 3]
adaptive-services {
  service-package layer-2;
}
mlfr-uni-nni-bundles 2; # Creates channelized LSQ interfaces/FRF.16 bundles.
[edit interfaces]
t1-0/0/0 {
  encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
    family mlfr-uni-nni {
      bundle lsq-1/3/0:1;
    }
  }
}
t1-0/0/1 {
  encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
    family mlfr-uni-nni {
      bundle lsq-1/3/0:1;
    }
  }
}
lsq-1/3/0:1 { # Bundle link consisting of t1-0/0/0 and t1-0/0/1
  per-unit-scheduler;
  encapsulation multilink-frame-relay-uni-nni;
  dce; # One end needs to be configured as DCE.
  mlfr-uni-nni-bundle-options {
    drop-timeout 180;
    fragment-threshold 64;
  }
}
```



```

hello-timer 180;
minimum-links 2;
mrru 3000;
link-layer-overhead 0.5;
}
unit 0 {
  dlci 26; # Each logical unit maps a single DLCI.
  family inet {
    address 10.2.3.4/24;
  }
}
unit 1 {
  dlci 42;
  family inet {
    address 10.20.30.40/24;
  }
}
unit 2 {
  dlci 69;
  family inet {
    address 10.20.30.40/24;
  }
}
[edit class-of-service]
scheduler-maps {
  sched-map-lsq0 {
    forwarding-class af scheduler af-scheduler-lsq0;
    forwarding-class be scheduler be-scheduler-lsq0;
    forwarding-class ef scheduler ef-scheduler-lsq0;
    forwarding-class nc scheduler nc-scheduler-lsq0;
  }
  sched-map-lsq1 {
    forwarding-class af scheduler af-scheduler-lsq1;
    forwarding-class be scheduler be-scheduler-lsq1;
    forwarding-class ef scheduler ef-scheduler-lsq1;
    forwarding-class nc scheduler nc-scheduler-lsq1;
  }
}
schedulers {
  af-scheduler-lsq0 {
    transmit-rate percent 60;
    buffer-size percent 60;
    priority low;
  }
  be-scheduler-lsq0 {
    transmit-rate percent 30;
    buffer-size percent 30;
    priority low;
  }
  ef-scheduler-lsq0 {
    transmit-rate percent 5;
    buffer-size percent 5;
    priority strict-high;
  }
  nc-scheduler-lsq0 {
    transmit-rate percent 5;

```

```
        buffer-size percent 5;
        priority high;
    }
    af-scheduler-lsq1 {
        transmit-rate percent 50;
        buffer-size percent 50;
        priority low;
    }
    be-scheduler-lsq1 {
        transmit-rate percent 30;
        buffer-size percent 30;
        priority low;
    }
    ef-scheduler-lsq1 {
        transmit-rate percent 15;
        buffer-size percent 15;
        priority strict-high;
    }
    nc-scheduler-lsq1 {
        transmit-rate percent 5;
        buffer-size percent 5;
        priority high;
    }
}
interfaces {
    lsq-1/3/0:1 { # MLFR FRF.16
        unit 0 {
            scheduler-map sched-map-lsq0;
        }
        unit 1 {
            scheduler-map sched-map-lsq1;
        }
    }
}
```

**Related
Documentation**

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using MLPPP on page 565](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.15 on page 576](#)
- [Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 589](#)
- [Link Services Configuration for Junos Interfaces](#)

Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.15

This example configures an NxT1 bundle using FRF.15 on a link services IQ interface. FRF.15 is similar to FRF.12, as described in [“Configuring LSQ Interfaces for Single Fractional T1 or El Interfaces Using FRF.12” on page 582](#). The difference is that FRF.15 supports multiple physical links in a bundle, whereas FRF.12 supports only one physical link per bundle. For the Junos OS implementation of FRF.15, you can configure one DLCI per physical link.



NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. This example refers to T1 interfaces, but the configuration for E1 interfaces is similar.

```
[edit interfaces]
lsq-1/3/0 {
  per-unit-scheduler;
  unit 0 {
    encapsulation multilink-frame-relay-end-to-end;
  }
}
unit 1 {
  encapsulation multilink-frame-relay-end-to-end;
}
# First physical link
t1-1/1/0:1 {
  encapsulation frame-relay;
  unit 0 {
    dlci 69;
    family mlfr-end-to-end {
      bundle lsq-1/3/0.0;
    }
  }
}
# Second physical link
t1-1/1/0:2 {
  encapsulation frame-relay;
  unit 0 {
    dlci 13;
    family mlfr-end-to-end {
      bundle lsq-1/3/0.0;
    }
  }
}
```

Related Documentation

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on page 565](#)
- [Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16 on page 571](#)
- [Link Services Configuration for Junos Interfaces](#)

Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI

When you configure a single fractional T1 interface, it is called a logical interface, because it can represent, for example, a routing adjacency.

The logical link services IQ interface represents the MLPPP bundle. Four queues are associated with the logical interface. A scheduler removes packets from the queues

according to a scheduling policy. Typically, you designate one queue to have strict priority, and the remaining queues are serviced in proportion to weights you configure.

To configure a single fractional T1 interface using MLPPP and LFI, you associate one DSO (fractional T1) interface with a link services IQ interface. To associate a fractional T1 interface with a link services IQ interface, include the **bundle** statement at the **[edit interfaces ds-fpc/pic/port:channel unit logical-unit-number family mlppp]** hierarchy level:

```
[edit interfaces ds-fpc/pic/port:channel unit logical-unit-number family mlppp]
bundle lsq-fpc/pic/port.logical-unit-number;
```



NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure the link services IQ interface properties, include the following statements at the **[edit interfaces lsq-fpc/pic/port unit logical-unit-number]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port unit logical-unit-number]
drop-timeout milliseconds;
encapsulation multilink-ppp;
fragment-threshold bytes;
link-layer-overhead percent;
minimum-links number;
mrru bytes;
short-sequence;
family inet {
    address address;
}
```

For MLPPP, assign a single scheduler map to the link services IQ (**lsq**) interface and to each constituent link. The default schedulers for M Series and T Series routers, which assign 95, 0, 0, and 5 percent bandwidth for the transmission rate and buffer size of queues 0, 1, 2, and 3, are not adequate when you configure LFI or multiclass traffic. Therefore, for MLPPP, you should configure a single scheduler with nonzero percent transmission rates and buffer sizes for queues 0 through 3, and assign this scheduler to the link services IQ (**lsq**) interface and to each constituent link and to each constituent link, as shown in [“Example: Configuring an LSQ Interface for a Fractional T1 Interface Using MLPPP and LFI” on page 580](#).



NOTE: For M320 and T Series routers, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

To configure and apply the scheduling policy, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
interfaces {
    ds-fpc/pic/port.channel {
```

```

        scheduler-map map-name;
    }
}
forwarding-classes {
    queue queue-number class-name;
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
schedulers {
    scheduler-name {
        buffer-size (percent percentage | remainder | temporal microseconds);
        priority priority-level;
        transmit-rate (rate | percent percentage | remainder) <exact>;
    }
}

```

For link services IQ interfaces, a strict-high-priority queue might starve all the other queues because traffic in a strict-high priority queue is transmitted before any other queue is serviced. This implementation is unlike the standard Junos CoS implementation in which a strict-high-priority queue receives infinite credits and does round-robin with high-priority queues, as described in the *Class of Service Feature Guide for Routing Devices*.

After the scheduler removes a packet from a queue, a certain action is taken. The action depends on whether the packet came from a multilink encapsulated queue (fragmented and sequenced) or a nonencapsulated queue (hashed with no fragmentation). Each queue can be designated as either multilink encapsulated or nonencapsulated, independently of the other. By default, traffic in all forwarding classes is multilink encapsulated. To configure packet fragmentation handling on a queue, include the **fragmentation-maps** statement at the **[edit class-of-service]** hierarchy level:

```

[edit class-of-service]
fragmentation-maps {
    map-name {
        forwarding-class class-name {
            fragment-threshold bytes;
            no-fragmentation;
        }
    }
}

```

If you require the queue to transmit small packets with low latency, configure the queue to be nonencapsulated by including the **no-fragmentation** statement. If you require the queue to transmit large packets with normal latency, configure the queue to be multilink encapsulated by including the **fragment-threshold** statement. If you require the queue to transmit large packets with low latency, we recommend using a faster link and configuring the queue to be nonencapsulated. For more information about fragmentation maps, see [“Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces” on page 527](#).

When a packet is removed from a multilink-encapsulated queue, it is fragmented. If the packet exceeds the minimum link MTU, or if a queue has a fragment threshold configured

at the **[edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]** hierarchy level, the software splits the packet into two or more fragments, which are assigned consecutive multilink sequence numbers.

If you do not include the **fragment-threshold** statement in the fragmentation map, the fragmentation threshold you set at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level is the default for all forwarding classes. If you do not set a maximum fragment size anywhere in the configuration, packets are fragmented if they exceed the smallest MTU of all the links in the bundle.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the maximum received reconstructed unit (MRRU) by including the **mrru** statement at the **[edit interfaces *lsq-fpc/pic/port* unit *logical-unit-number*]** hierarchy level. The MRRU is similar to the MTU, but is specific to link services interfaces. By default the MRRU size is 1500 bytes, and you can configure it to be from 1500 through 4500 bytes. For more information, see [“Configuring MRRU on Multilink and Link Services Logical Interfaces” on page 734](#).

When a packet is removed from a multilink-encapsulated queue, the software gives the packet an MLPPP header. The MLPPP header contains a sequence number field, which is filled with the next available sequence number from a counter. The software then places the packet on the fractional T1 link. Traffic from another queue might be interleaved between two fragments of the packet.

When a packet is removed from a nonencapsulated queue, it is transmitted with a plain PPP header. The packet is then placed on the fractional T1 link as soon as possible. If necessary, the packet is placed between the fragments of a packet from another queue.

The fractional T1 interface links to another router, which can be from Juniper Networks or another vendor. The router at the far end gathers packets from the fractional T1 link. If a packet has an MLPPP header, the software assumes the packet is a fragment of a larger packet, and the fragment number field is used to reassemble the larger packet. If the packet has a plain PPP header, the software accepts the packet in the order in which it arrives, and the software makes no attempt to reassemble or reorder the packet.

Example: Configuring an LSQ Interface for a Fractional T1 Interface Using MLPPP and LFI

Configure a single fractional T1 logical interface:

```
[edit interfaces]
lsq-0/2/0 {
  per-unit-scheduler;
  unit 0 {
    encapsulation multilink-ppp;
    link-layer-overhead 0.5;
    family inet {
      address 10.40.1.1/30;
    }
  }
}
ct3-1/0/0 {
  partition 1 interface-type ct1;
}
```

```

ct1-1/0/0:1 {
    partition 1 timeslots 1-2 interface-type ds;
}
ds-1/0/0:1:1 {
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle lsq-0/2/0.0;
        }
    }
}
[edit class-of-service]
interfaces {
    ds-1/0/0:1:1 { # multilink PPP constituent link
        unit 0 {
            scheduler-map sched-map1;
        }
    }
}
forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
}
scheduler-maps {
    sched-map1 {
        forwarding-class af scheduler af-scheduler;
        forwarding-class be scheduler be-scheduler;
        forwarding-class ef scheduler ef-scheduler;
        forwarding-class nc scheduler nc-scheduler;
    }
}
schedulers {
    af-scheduler {
        transmit-rate percent 20;
        buffer-size percent 20;
        priority low;
    }
    be-scheduler {
        transmit-rate percent 20;
        buffer-size percent 20;
        priority low;
    }
    ef-scheduler {
        transmit-rate percent 50;
        buffer-size percent 50;
        priority strict-high; # voice queue
    }
    nc-scheduler {
        transmit-rate percent 10;
        buffer-size percent 10;
        priority high;
    }
}
fragmentation-maps {
    fragmap-1 {

```

```

        forwarding-class be {
            fragment-threshold 180;
        }
        forwarding-class ef {
            fragment-threshold 100;
        }
    }
}
[edit interfaces]
lsq-0/2/0 {
    unit 0 {
        fragmentation-map fragmap-1;
    }
}

```

Related Documentation

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12 on page 582](#)
- [Link Services Configuration for Junos Interfaces](#)

Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12

To configure a single fractional T1 interface using FRF.16, you associate a DS0 interface with a link services IQ (lsq) interface. When you configure a single fractional T1, the fractional T1 carries a potentially large number of Frame Relay PVCs identified by their DLCIs. Each DLCI is called a logical interface, because it can represent, for example, a routing adjacency. To associate the DS0 interface with a link services IQ interface, include the **bundle** statement at the **[edit interfaces ds-fpc/pic/port:channel unit logical-unit-number family mlfr-end-to-end]** hierarchy level:

```
[edit interfaces ds-fpc/pic/port:channel unit logical-unit-number family mlfr-end-to-end]
bundle lsq-fpc/pic/port.logical-unit-number;
```



NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure the link services IQ interface properties, include the following statements at the **[edit interfaces lsq-fpc/pic/port unit logical-unit-number]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port unit logical-unit-number]
drop-timeout milliseconds;
encapsulation multilink-frame-relay-end-to-end;
fragment-threshold bytes;
link-layer-overhead percent;
minimum-links number;
mrru bytes;
short-sequence;
family inet {
    address address;
```



```
}
```

The logical link services IQ interface represents the FRF.12 bundle. Four queues are associated with each logical interface. A scheduler removes packets from the queues according to a scheduling policy. Typically, you designate one queue to have strict priority, and the remaining queues are serviced in proportion to weights you configure.

For FRF.12, assign a single scheduler map to the link services IQ interface (**lsq**) and to each constituent link. For M Series and T Series routers, the default schedulers, which assign 95, 0, 0, and 5 percent bandwidth for the transmission rate and buffer size of queues 0, 1, 2, and 3, are not adequate when you configure LFI or multiclass traffic. Therefore, for FRF.12, you should configure schedulers with nonzero percent transmission rates and buffer sizes for queues 0 through 3, and assign them to the link services IQ interface (**lsq**) and to each constituent link, as shown in [“Examples: Configuring an LSQ Interface for a Fractional T1 Interface Using FRF.12” on page 585](#).



NOTE: For M320 and T Series routers, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

To configure and apply the scheduling policy, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
interfaces {
  ds-fpc/pic/port.channel {
    scheduler-map map-name;
  }
}
forwarding-classes {
  queue queue-number class-name;
}
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
schedulers {
  scheduler-name {
    buffer-size (percent percentage | remainder | temporal microseconds);
    priority priority-level;
    transmit-rate (rate | percent percentage | remainder) <exact>;
  }
}
```

For link services IQ interfaces, a strict-high-priority queue might starve the other three queues because traffic in a strict-high-priority queue is transmitted before any other queue is serviced. This implementation is unlike the standard Junos CoS implementation in which a strict-high-priority queue does round-robin with high-priority queues, as described in the *Class of Service Feature Guide for Routing Devices*.

After the scheduler removes a packet from a queue, a certain action is taken. The action depends on whether the packet came from a multilink encapsulated queue (fragmented and sequenced) or a nonencapsulated queue (hashed with no fragmentation). Each queue can be designated as either multilink encapsulated or nonencapsulated, independently of the other. By default, traffic in all forwarding classes is multilink encapsulated. To configure packet fragmentation handling on a queue, include the **fragmentation-maps** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      fragment-threshold bytes;
      no-fragmentation;
    }
  }
}
```

If you require the queue to transmit small packets with low latency, configure the queue to be nonencapsulated by including the **no-fragmentation** statement. If you require the queue to transmit large packets with normal latency, configure the queue to be multilink encapsulated by including the **fragment-threshold** statement. If you require the queue to transmit large packets with low latency, we recommend using a faster link and configuring the queue to be nonencapsulated. For more information about fragmentation maps, see [“Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces” on page 527](#).

When a packet is removed from a multilink-encapsulated queue, it is fragmented. If the packet exceeds the minimum link MTU, or if a queue has a fragment threshold configured at the **[edit class-of-service fragmentation-maps map-name forwarding-class class-name]** hierarchy level, the software splits the packet into two or more fragments, which are assigned consecutive multilink sequence numbers.

If you do not include the **fragment-threshold** statement in the fragmentation map, the fragmentation threshold you set at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level is the default for all forwarding classes. If you do not set a maximum fragment size anywhere in the configuration, packets are fragmented if they exceed the smallest MTU of all the links in the bundle.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the maximum received reconstructed unit (MRRU) by including the **mrru** statement at the **[edit interfaces lsq-fpc/pic/port unit logical-unit-number]** hierarchy level. The MRRU is similar to the MTU but is specific to link services interfaces. By default, the MRRU size is 1500 bytes, and you can configure it to be from 1500 through 4500 bytes. For more information, see [“Configuring MRRU on Multilink and Link Services Logical Interfaces” on page 734](#).

When a packet is removed from a multilink-encapsulated queue, the software gives the packet an FRF.12 header. The FRF.12 header contains a sequence number field, which is filled with the next available sequence number from a counter. The software then places the packet on the fractional T1 link. Traffic from another queue might be interleaved between two fragments of the packet.

When a packet is removed from a nonencapsulated queue, it is transmitted with a plain Frame Relay header. The packet is then placed on the fractional T1 link as soon as possible. If necessary, the packet is placed between the fragments of a packet from another queue.

The fractional T1 interface links to another router, which can be from Juniper Networks or another vendor. The router at the far end gathers packets from the fractional T1 link. If a packet has an FRF.12 header, the software assumes the packet is a fragment of a larger packet, and the fragment number field is used to reassemble the larger packet. If the packet has a plain Frame Relay header, the software accepts the packet in the order in which it arrives, and the software makes no attempt to reassemble or reorder the packet.

A whole packet from a nonencapsulated queue can be placed between fragments of a multilink-encapsulated queue. However, fragments from one multilink-encapsulated queue cannot be interleaved with fragments from another multilink-encapsulated queue. This is the intent of the specification FRF.12, *Frame Relay Fragmentation Implementation Agreement*. If fragments from two different queues were interleaved, the header fields might not have enough information to separate the fragments.

Examples: Configuring an LSQ Interface for a Fractional T1 Interface Using FRF.12

FRF.12 with Fragmentation and Without LFI

This example shows a 128 KB DS0 interface. There is one traffic stream on **ge-0/0/0**, which is classified into queue 0 (**be**). Packets are fragmented in the link services IQ (**lsq-**) interface according to the threshold configured in the fragmentation map.

```
[edit chassis]
fpc 0 {
  pic 3 {
    adaptive-services {
      service-package layer-2;
    }
  }
}
[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 20.1.1.1/24 {
        arp 20.1.1.2 mac 00.90.1b.12.34.56;
      }
    }
  }
}
cel-0/2/0 {
  partition 1 timeslots 1-2 interface-type ds;
}
ds-0/2/0:1 {
  no-keepalives;
  dce;
  encapsulation frame-relay;
```

```
    unit 0 {
        dlci 100;
        family mfr-end-to-end {
            bundle lsq-0/3/0.0;
        }
    }
}
lsq-0/3/0 {
    per-unit-scheduler;
    unit 0 {
        encapsulation multilink-frame-relay-end-to-end;
        family inet {
            address 10.200.0.78/30;
        }
    }
}
fxp0 {
    unit 0 {
        family inet {
            address 172.16.1.162/24;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.1/32;
        }
    }
}
[edit class-of-service]
forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
}
interfaces {
    lsq-0/3/0 {
        unit 0 {
            fragmentation-map map1;
        }
    }
}
fragmentation-maps {
    map1 {
        forwarding-class {
            be {
                fragment-threshold 160;
            }
        }
    }
}
}
```

**FRF.12 with
Fragmentation and LFI**

This example shows a 512 KB DSO bundle and four traffic streams on **ge-0/0/0** that are classified into four queues. The fragment size is 160 for queue 0, queue 1, and queue 2. The voice stream on queue 3 has LFI configured.

```
[edit chassis]
fpc 0 {
  pic 3 {
    adaptive-services {
      service-package layer-2;
    }
  }
}
[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 20.1.1.1/24 {
        arp 20.1.1.2 mac 00.90.1b.12.34.56;
      }
    }
  }
  ce1-0/2/0 {
    partition 1 timeslots 1-8 interface-type ds;
  }
  ds-0/2/0:1 {
    no-keepalives;
    dce;
    encapsulation frame-relay;
    unit 0 {
      dlci 100;
      family mlfr-end-to-end {
        bundle lsq-0/3/0.0;
      }
    }
  }
}
lsq-0/3/0 {
  per-unit-scheduler;
  unit 0 {
    encapsulation multilink-frame-relay-end-to-end;
    family inet {
      address 10.200.0.78/30;
    }
  }
}
[edit class-of-service]
classifiers {
  inet-precedence ge-interface-classifier {
    forwarding-class be {
      loss-priority low code-points 000;
    }
    forwarding-class ef {
      loss-priority low code-points 010;
    }
    forwarding-class af {
      loss-priority low code-points 100;
    }
  }
}
```

```
    }
    forwarding-class nc {
        loss-priority low code-points 110;
    }
}
forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
}
interfaces {
    lsq-0/3/0 {
        unit 0 {
            scheduler-map sched2;
            fragmentation-map map2;
        }
    }
    ds-0/2/0:1 {
        scheduler-map link-map2;
    }
    ge-0/0/0 {
        unit 0 {
            classifiers {
                inet-precedence ge-interface-classifier;
            }
        }
    }
}
scheduler-maps {
    sched2 {
        forwarding-class be scheduler economy;
        forwarding-class ef scheduler business;
        forwarding-class af scheduler stream;
        forwarding-class nc scheduler voice;
    }
    link-map2 {
        forwarding-class be scheduler link-economy;
        forwarding-class ef scheduler link-business;
        forwarding-class af scheduler link-stream;
        forwarding-class nc scheduler link-voice;
    }
}
fragmentation-maps {
    map2 {
        forwarding-class {
            be {
                fragment-threshold 160;
            }
            ef {
                fragment-threshold 160;
            }
            af {
                fragment-threshold 160;
            }
        }
    }
}
```

```

        nc {
            no-fragmentation;
        }
    }
}
schedulers {
    economy {
        transmit-rate percent 26;
        buffer-size percent 26;
    }
    business {
        transmit-rate percent 26;
        buffer-size percent 26;
    }
    stream {
        transmit-rate percent 35;
        buffer-size percent 35;
    }
    voice {
        transmit-rate percent 13;
        buffer-size percent 13;
    }
    link-economy {
        transmit-rate percent 26;
        buffer-size percent 26;
    }
    link-business {
        transmit-rate percent 26;
        buffer-size percent 26;
    }
    link-stream {
        transmit-rate percent 35;
        buffer-size percent 35;
    }
    link-voice {
        transmit-rate percent 13;
        buffer-size percent 13;
    }
}
}
}

```

**Related
Documentation**

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI on page 577](#)
- *Link Services Configuration for Junos Interfaces*

Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP

This example bundles a single T3 interface on a link services IQ interface with MLPPP encapsulation. Binding a single T3 interface to a multilink bundle allows you to configure compressed RTP (CRTP) on the T3 interface.

This scenario applies to MLPPP bundles only. The Junos OS does not currently support CRTP over Frame Relay. For more information, see [“Configuring Services Interfaces for Voice Services” on page 622](#).

There is no need to configure LFI at DS3 speeds, because the packet serialization delay is negligible.

```
[edit interfaces]
t3-0/0/0 {
  unit 0 {
    family mlppp {
      bundle lsq-1/3/0.1;
    }
  }
}
lsq-1/3/0.1 {
  encapsulation multilink-ppp;
}
compression {
  rtp {
    # cRTP parameters go here
    #
    port minimum 2000 maximum 64009;
  }
}
```

This configuration uses a default fragmentation map, which results in all forwarding classes (queues) being sent out with a multilink header.

To eliminate multilink headers, you can configure a fragmentation map in which all queues have the **no-fragmentation** statement at the **[edit class-of-service fragmentation-maps map-name forwarding-class class-name]** hierarchy level, and attach the fragmentation map to the **lsq-1/3/0.1** interface, as shown here:

```
[edit class-of-service]
fragmentation-maps {
  fragmap {
    forwarding-class {
      be {
        no-fragmentation;
      }
      af {
        no-fragmentation;
      }
      ef {
        no-fragmentation;
      }
      nc {
        no-fragmentation;
      }
    }
  }
}
interfaces {
  lsq-1/3/0.1 {
```



```

        fragmentation-map fragmap;
    }
}

```

Related Documentation

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on page 565](#)
- [Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI on page 577](#)
- [Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP on page 593](#)
- [Link Services Configuration for Junos Interfaces](#)

Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12

This example configures a clear-channel T3 or OC3 interface with multiple logical interfaces (DLCIs) on the link. In this scenario, each DLCI represents a customer. DLCIs are shaped at the egress PIC to a particular speed ($N \times \text{DSO}$). This allows you to configure LFI using FRF.12 End-to-End Protocol on Frame Relay DLCIs.

To do this, first configure logical interfaces (DLCIs) on the physical interface. Then bundle the DLCIs, so that there is only one DLCI per bundle.

The physical interface must be capable of per-DLCI scheduling, which allows you to attach shaping rates to each DLCI. For more information, see the *Junos OS Network Interfaces Library for Routing Devices*.

To prevent fragment drops at the egress PIC, you must assign a shaping rate to the link services IQ logical interfaces and to the egress DLCIs. Shaping rates on DLCIs specify how much bandwidth is available for each DLCI. The shaping rate on link services IQ interfaces should match the shaping rate assigned to the DLCI that is associated with the bundle.

Egress interfaces also must have a scheduler map attached. The queue that carries voice should be strict-high-priority, while all other queues should be low-priority. This makes LFI possible.

This example shows voice traffic in the `ef` queue. The voice traffic is interleaved with bulk data. Alternatively, you can use multiclass MLPPP to carry multiple classes of traffic in different multilink classes, as described in [“Configuring Multiclass MLPPP on LSQ Interfaces” on page 562](#).

```

[edit interfaces]
t3-0/0/0 {
    per-unit-scheduler;
    encapsulation frame-relay;
    unit 0 {
        dlc1 69;
        family mlfr-end-to-end {
            bundle lsq-1/3/0.0;
        }
    }
}

```

```
    unit 1 {
        dlc1 42;
        family mfr-end-to-end {
            bundle lsq-1/3/0.1;
        }
    }
}
lsq-1/3/0 {
    unit 0 {
        encapsulation multilink-frame-relay-end-to-end;
    }
    fragment-threshold 320; # Multilink packets must be fragmented
}
unit 1 {
    encapsulation multilink-frame-relay-end-to-end;
}
fragment-threshold 160;
[edit class-of-service]
scheduler-maps {
    sched { # Scheduling parameters that apply to bundles on AS or Multiservices PICs.
        ...
    }
    pic-sched {
        # Scheduling parameters for egress DLCIs.
        # The voice queue should be strict-high priority.
        # All other queues should be low priority.
        ...
    }
}
fragmentation-maps {
    fragmap {
        forwarding-class {
            ef {
                no-fragmentation;
            }
            # Voice is carried in the ef queue.
            # It is interleaved with bulk data.
        }
    }
}
}
interfaces {
    t3-0/0/0 {
        unit 0 {
            shaping-rate 512k;
            scheduler-map pic-sched;
        }
        unit 1 {
            shaping-rate 128k;
            scheduler-map pic-sched;
        }
    }
}
lsq-1/3/0 { # Assign fragmentation and scheduling to LSQ interfaces.
    unit 0 {
        shaping-rate 512k;
        scheduler-map sched;
        fragmentation-map fragmap;
    }
}
```

```

unit 1 {
    shaping-rate 128k;
    scheduler-map sched;
    fragmentation-map fragmap;
}

```

For more information about how FRF.12 works with links services IQ interfaces, see [“Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12” on page 582.](#)

**Related
Documentation**

- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
- [Link Services Configuration for Junos Interfaces](#)
- [Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 589](#)

Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP

This example configures an ATM2 IQ interface with MLPPP bundled with link services IQ interfaces. This allows you to configure LFI on ATM virtual circuits.

For this type of configuration, the ATM2 IQ interface must have LLC encapsulation.

The following ATM PICs are supported in this scenario:

- 2-port OC-3/STM1 ATM2 IQ
- 4-port DS3 ATM2 IQ

Virtual circuit multiplexed PPP over AAL5 is not supported. Frame Relay is not supported. Bundling of multiple ATM VCs into a single logical interface is not supported.

Unlike DS3 and OC3 interfaces, there is no need to create a separate scheduler map for the ATM PIC. For ATM, you define CoS components at the **[edit interfaces at-fpc/pic/port atm-options]** hierarchy level, as described in the *Junos OS Network Interfaces Library for Routing Devices*.



NOTE: Do not configure RED profiles on ATM logical interfaces that are bundled. Drops do not occur at the ATM interface.

In this example, two ATM VCs are configured and bundled into two link services IQ bundles. A fragmentation map is used to interleave voice traffic with other multilink traffic. Because MLPPP is used, each link services IQ bundle can be configured for CRTP.

```

[edit interfaces]
at-1/2/0 {
    atm-options {
        vpi 0;
        pic-type atm2;
    }
}

```

```
    unit 0 {
        vci 0.69;
        encapsulation atm-mlppp-llc;
        family mlppp {
            bundle lsq-1/3/0.10;
        }
    }
    unit 1 {
        vci 0.42;
        encapsulation atm-mlppp-llc;
        family mlppp {
            bundle lsq-1/3/0.11;
        }
    }
}
lsq-1/3/0 {
    unit 10 {
        encapsulation multilink-ppp;
    }
    # Large packets must be fragmented.
    # You can specify fragmentation for each forwarding class.
    fragment-threshold 320;
    compression {
        rtp {
            port minimum 2000 maximum 64009;
        }
    }
}
unit 11 {
    encapsulation multilink-ppp;
}
fragment-threshold 160;
[edit class-of-service]
scheduler-maps {
    sched { # Scheduling parameters that apply to LSQ bundles on AS or Multiservices PICs.
        ...
    }
}
fragmentation-maps {
    fragmap {
        forwarding-class {
            ef {
                no-fragmentation;
            }
        }
    }
}
}
interfaces { # Assign fragmentation and scheduling parameters to LSQ interfaces.
lsq-1/3/0 {
    unit 0 {
        shaping-rate 512k;
        scheduler-map sched;
        fragmentation-map fragmap;
    }
    unit 1 {
        shaping-rate 128k;
        scheduler-map sched;
    }
}
```

```
        fragmentation-map fragmap;  
    }  
}
```

- Related Documentation**
- [Layer 2 Service Package Capabilities and Interfaces on page 543](#)
 - *Link Services Configuration for Junos Interfaces*

PART 10

Enabling Load Balancing and High Availability Using Multiservices Interfaces

- [Enabling Load Balancing and High Availability Using Multiservices Interfaces on page 599](#)

CHAPTER 41

Enabling Load Balancing and High Availability Using Multiservices Interfaces

- [Understanding Aggregated Multiservices Interfaces on page 599](#)
- [Configuring Load Balancing on AMS Infrastructure on page 605](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 608](#)
- [Example: Configuring Next-Hop Style Services on an Aggregated Multiservices Interface on page 613](#)
- [Example: Configuring Static Source Translation on AMS Infrastructure on page 616](#)

Understanding Aggregated Multiservices Interfaces

This topic contains the following sections:

- [Aggregated Multiservices Interface on page 599](#)
- [IPv6 Traffic on AMS Interfaces Overview on page 603](#)
- [Member Failure Options and High Availability Settings on page 604](#)

Aggregated Multiservices Interface

In Junos OS, you can combine multiple services interfaces to create a bundle of interfaces that can function as a single interface. Such a bundle of interfaces is known as an aggregated multiservices interface (AMS), and is denoted as `amsN` in the configuration, where *N* is a unique number that identifies an AMS interface (for example, `ams0`).

AMS configuration provides higher scalability, improved performance, and better failover and load-balancing options.

The current service set configuration model in Junos OS supports only one service PIC per service set. All services provisioned using a service set must be handled by the only one service PIC associated with that service set. AMS configuration enables you to address this limitation by associating an AMS bundle with a service set. An AMS bundle can have up to eight services PICs as member interfaces and can distribute services among the member interfaces. This allows you to have multiple service interfaces to handle services configured in one service set.

Member interfaces are identified as **mams** in the configuration. The **chassisd** process in routers that support AMS configuration creates a **mams** entry for every multiservices interface on the router.

When you configure **services-options** at the **ams** interface level, the options apply to all member interfaces (**mams**) for the **ams** interface.

The options also apply to service sets configured on **ms-** interfaces corresponding to the **ams** interface's member interfaces. All settings are per PIC. For example, session-limit applies per member and not at an aggregate level.



NOTE: You cannot configure **services-options** at both the **ams** (aggregate) and member-interface level. If **services-options** is configured on **ms-x/y/z**, it also applies to service sets on **mams-x/y/z**.

Some different options are available under **services-options** on **ms-** interfaces, so the user should ensure that differences in configuration is meaningful [For example, timeouts should be similar but, say, **syslog-prefix/ip-address** can be different.



NOTE: Per-member drop of traffic and per-member next-hop configuration is required for NAT64. For NAPT-44, this per-member specification allows arbitrary hash-keys and therefore this setting enables better load-balancing options. The main purpose is to allow dynamic NAT operations to be performed. For NAT64, NAPT44, and dynamic NAT44, it is not possible to determine which member allocates the dynamic NAT address. To ensure that reverse flow packets arrive at the same member as the forward flow packets, pool-address-based routes are used to steer reverse flow packets.



NOTE: Until Junos OS Release 13.3, for every media logical interface on which services were configured (interface style services), a logical interface alias was internally created. This interface alias stores the topology chains for features that are performed on the logical interface after an input service was processed to avoid packet loops in the system. With interface aliases, the maximum number of logical interfaces supported with services was reduced to half the supported maximum number because each logical interface consumed two entries, namely, one for the interface itself and the other for the interface alias.

Starting with Junos OS Release 14.1R4, input interface aliases are not created for MS-MPCs and MS-MICs. As a result, the maximum number of logical interfaces that are supported with services PICs is equal to the maximum number supported on the system. After input service processing by MS-MPCs and MS-MICs, the services PIC sends the packet to the Packet Forwarding Engine on the multiservices (ms-) logical interface where the corresponding service is configured. Post-services are not supported on MS-MPCs and MS-MICs in Junos OS Release 13.2 and later.



NOTE: You cannot include MS-DPCs or other multiservices PICs in an AMS configuration that contains MS-MICs or MS-MPCs as member interfaces.

By default, the traffic distribution over the member interfaces of an AMS interface happens in a round-robin fashion. You can also configure the following hash key values to regulate the traffic distribution: **source-ip**, **destination-ip**, **iif** (incoming interface), **oif** (outgoing interface), and **protocol**. For services that require traffic symmetry, you must configure symmetrical hashing. Symmetrical hashing configuration ensures that both forward and reverse traffic are routed through the same member interface.



NOTE:



NOTE: With basic NAT44, load balancing on AMS interfaces of MS-MICs and MS-MPCs does not work properly if the ingress hash key is source IP address and the egress hash key is destination IP address.

If the service set is applied on the Gigabit Ethernet or 10-Gigabit Ethernet interface that functions as the NAT inside interface, then the hash keys used for load balancing might be configured in such a way that the ingress key is set as destination IP address and the egress key is set as source IP address. Because the source IP address undergoes NAT processing, it is not available for hashing the traffic in the reverse direction. Therefore, load-balancing does not happen on the same IP address and forward and reverse traffic do not map to the same PIC. With the hash keys reversed, load balancing occurs correctly.

With next-hop services, for forward traffic, the ingress-key on the inside-interface load-balances traffic, and for reverse traffic, the ingress-key on the outside-interface load-balances traffic or per-member-next-hops steer reverse traffic. With interface-style services, the ingress-key load-balances forward traffic and the egress-key load-balances forward traffic or per-member-next-hops steer reverse traffic. Forward traffic is traffic entering from the inner side of a service-set and reverse traffic is traffic entering from the outer side of a service-set. The forward key is the hash key used for the forward direction of traffic and the reverse key is the hash key used for the reverse direction of traffic (depends on whether it relates to interface-services or next-hop-services style.)

With stateful firewalls, you can configure the following combinations of forward and reverse keys for load balancing. In the following combinations presented for hash keys, FOR-KEY refers to the forward key, REV-KEY denotes the reverse key, SIP signifies source IP address, DIP signifies destination IP address, and PROTO refers to protocol such as IP.

- FOR-KEY: SIP, REV-KEY: DIP
- FOR-KEY: SIP,PROTO REV-KEY: DIP, PROTO
- FOR-KEY: DIP, REV-KEY: SIP
- FOR-KEY: DIP,PROTO REV-KEY: SIP, PROTO
- FOR-KEY: SIP,DIP REV-KEY: SIP, DIP
- FOR-KEY: SIP,DIP,PROTO REV-KEY: SIP, DIP,PROTO

With static NAT configured as basic NAT44 or destination NAT44, and with stateful firewall configured or not, if the forward direction of traffic must undergo NAT processing, configure the hash keys as follows:

- FOR-KEY: DIP, REV-KEY: SIP
- FOR-KEY: DIP,PROTO REV-KEY: SIP, PROTO

If the reverse direction of traffic must undergo NAT processing, configure the hash keys as follows:

- FOR-KEY: SIP, REV-KEY: DIP
- FOR-KEY: SIP,PROTO REV-KEY: DIP, PROTO

With dynamic NAT configured, and with stateful firewall configured or not, only the forward direction traffic can undergo NAT. The forward hash key can be any combination of SIP, DIP, and protocol, and the reverse hash key is ignored.



NOTE:



NOTE: Junos OS AMS configuration supports IPv4 and IPv6 traffic.

IPv6 Traffic on AMS Interfaces Overview

Starting in Junos OS release 14.2R1, you can use AMS interfaces for IPv6 traffic. To configure IPv6 support for an AMS interface, include the **family inet6** statement at the **[edit interfaces *ams-interface-name* unit 1]** hierarchy level. When **family inet** and **family inet6** are set for an AMS interface sub-unit, the **hash-keys** configured at the **[edit services *service-set-name* load-balancing-options]** hierarchy level apply to both the IPv4 and IPv6 flows.

With the support for transmission of IPv6 packets on AMS interfaces, the redistribution action that occurs, when an AMS member interface goes down, has been enhanced. When a member interface of an AMS bundle fails, traffic destined to the failed member is redistributed among the remaining active members. The traffic (flows or sessions) traversing through the existing active members is unaffected. If M members are currently active, the expected result is that only about 1/M fraction of the traffic (flows/sessions) is impacted because that amount of traffic is shifted from the failed member to remain active members. When the failed member interface comes back online, only a fraction of the traffic is redistributed to the new member. If N members are currently active, the expected result is that only about 1/(N+1) fraction of the traffic (flows/sessions) is impacted because that amount of traffic will be moved to the new restored member. The aforementioned values of 1/M and 1/(N+1) assume that the flows are uniformly distributed among members. Because a packet-hash is used to load-balance and because traffic usually contains a typical random combination of IP addresses (or any other fields that are used as load-balancing keys), this assumptions hold good.

Similar to IPv4 traffic, for IPv6 packets, an AMS bundle must contain members of only one service PIC type. You cannot combine MS-DPC(XLR), MS-MIC(XLP) and MS-MPC(XLP) line cards to be of the same AMS bundle. Such service PICs can however be members of separate AMS bundles on the same router (for example, two MS-MICs in *ams0*, and two MS-MPC PICs in *ams1*). The number of flows distributed, in an ideal

environment, can be $1/N$ in a best-case scenario when the N th member goes up or down. However, this assumption considers that the hash-keys load-balance the real or dynamic traffic. For example, consider a real-world deployment where member A is serving only one flow, whereas member B is serving ten flows. If member B goes down, then the number of flows disrupted is 10/11. The NAT pool-split behavior is designed to utilize the benefits of the rehash-minimization feature. The splitting of a NAT pool is performed for dynamic NAT scenarios (dynamic NAT, NAT64, and NATPT44).

If the "original" and "redistributed" flows are defined as follows:

- Member-original-flows—The traffic mapped to a member when all members are up.
- Member-redistributed-flows—The additional traffic mapped to a member, when some other member fails. These traffic flows might need to be rebalanced when member interfaces come up and go down.

With the preceding definitions of the original and redistributed flows for member interfaces, the following observations apply:

- The member-original-flows of a member stay intact as long as that member is up. Such flows are not impacted when other members move between the up and down states.
- The member-redistributed-flows of a member can change when other members go up or down. This change of flows occurs because these additional flows need to be rebalanced among all active members. Therefore, the member-redistributed-flow can vary a lot based on other members going down or up. Although it might seem that when a member goes down, the flows on active-members are preserved, and that when a member goes up, flows on active-members are not preserved in an effective way, this behavior is only because of static or hash-based rebalancing of traffic among active members.

The rehash-minimization feature handles the operational changes in a member interface status only (such as member offline or member OS reset). It does not handle changes in configuration. For example, addition or deletion, or activation and deactivation, of member interfaces at the **[edit interfaces amsN load-balancing-options member-interface mams-a/b/0]** hierarchy level requires the member PICs to be bounced. Twice NAT or hairpinning is not supported, similar to IPv4 support for AMS interfaces.

Member Failure Options and High Availability Settings

Because multiple service interfaces are configured as part of an AMS bundle, AMS configuration also provides for failover and high availability support. You can either configure one of the member interfaces as a backup interface that becomes active when any one of the other member interfaces goes down, or configure the AMS in such a way that when one of the member interfaces goes down, the traffic assigned to that interface is shared across the active interfaces.

The **member-failure-options** configuration statement enables you to configure how to handle traffic when a member interface fails. One option is to redistribute the traffic immediately among the other member interfaces. However, redistribution of traffic

involves recalculating the hash tags, and might cause some disruption in traffic on all the member interfaces.

The other option is to configure the AMS to drop all traffic that is assigned to the failed member interface. With this you can optionally configure an interval, **rejoin-timeout**, for the AMS to wait for the failed interface to come back online after which the AMS can redistribute the traffic among other member interfaces. If the failed member interface comes back online before the configured wait time, traffic continues unaffected on all member interfaces, including the interface that has come back online and resumed the operations.

You can also control the rejoining of the failed interface when it comes back online. If you do not include the **enable-rejoin** statement in the **member-failure-options** configuration, the failed interface is not allowed to rejoin the AMS when it comes back online. In such cases, you can manually rejoin that to the AMS by executing the **request interfaces revert interface-name** operational mode command.

The **rejoin-timeout** and **enable-rejoin** statements enable you to minimize traffic disruptions when member interfaces flap.



NOTE: When **member-failure-options** are not configured, the default behavior is to drop member traffic with a rejoin timeout of 120 seconds.

The **high-availability-options** configuration enables you to designate one of the member interfaces as a backup interface. The backup interface does not participate in routing operations as long as it remains a backup interface. When a member interface fails, the backup interface handles the traffic assigned to the failed interface. When the failed interface comes back online, it becomes the new backup interface.

When both **member-failure-options** and **high-availability-options** are configured for an AMS, the **high-availability-options** configuration takes precedence over the **member-failure-options** configuration. If a second failure occurs before the failed interface comes back online to be the new backup, the **member-failure-options** configuration comes into effect.

**Related
Documentation**

- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 608](#)
- [Example: Configuring Next-Hop Style Services on an Aggregated Multiservices Interface on page 613](#)

Configuring Load Balancing on AMS Infrastructure

Configuring load balancing requires an aggregated Multiservices (AMS) system. AMS involves grouping several Multiservices PICs together. An AMS configuration eliminates the need for separate routers within a system. The primary benefit of having an AMS configuration is the ability to support load balancing of traffic across multiple services PICs.



NOTE: AMS is supported only on Mobility Gateway (MBG) with the MBG MS-DPC. AMS is not supported with JUNOS services like NAT, FW, IPsec, DAA, HCM on the current MS-DPC.

Starting with Junos OS 11.4, high availability (HA) is supported on AMS infrastructure on all MX Series 3D Universal Edge routers. AMS has several benefits:

- Support for configuring behavior if a Multiservices PIC that is part of the AMS configuration fails
- Support for specifying hash keys for each service set in either direction
- Support for adding routes to individual PICs within the AMS system

Configuring AMS Infrastructure

AMS supports load balancing across multiple service sets. All ingress or egress traffic for a service set can be load balanced across different services PICs. To enable load balancing, you have to configure an aggregate interface with existing services interfaces.

To configure failure behavior in AMS, include the **member-failure-options** statement:

```
[edit interfaces ams1]
load-balancing-options {
  member-failure-options {
    drop-member-traffic {
      rejoin-timeout rejoin-timeout;
    }
    redistribute-all-traffic {
      enable-rejoin;
    }
  }
}
```

If a PIC fails, the traffic to the failed PIC can be configured to be redistributed by using the **redistribute-all-traffic** statement at the **[edit interfaces *interface-name* load-balancing-options member-failure-options]** hierarchy level. If the **drop-member-traffic** statement is used, all traffic to the failed PIC is dropped. Both options are mutually exclusive.



NOTE: If **member-failure-options** is not explicitly configured, the default behavior is to drop member traffic with a rejoin timeout of 120 seconds.

Only **mams-** interfaces (services interfaces that are part of AMS) can be aggregated. After an AMS interface has been configured, the constituent **mams-** interfaces cannot be individually configured. A **mams-** interface cannot be used as an **rms** interface. AMS supports IPv4 (family inet) and IPv6 (family inet6). It is not possible to configure addresses on an AMS interface. Network Address Translation (NAT) is the only application that runs on AMS infrastructure at this time.



NOTE: Unit 0 on an AMS interface cannot be configured.

To support multiple applications and different types of translation, AMS infrastructure supports configuring hashing for each service set. The hash keys can be configured separately for ingress and egress. The default configuration uses source IP, destination IP, and the protocol for hashing; incoming-interface for ingress and outgoing-interface for egress are also available.

Configuring High Availability

In an AMS system configured with high availability, a designated Multiservices PIC acts as a backup for other active PICs that are part of the AMS system. Presently, only N:1 backup for high availability is supported; only one PIC is available as backup for all other active PICs. High availability for load balancing is configured by adding the **high-availability-options** statement at the **[edit interfaces *interface-name* load-balancing-options]** hierarchy level.

To configure high availability, include the **high-availability-options** statement:

```
[edit interfaces ams1]
load-balancing-options {
  high-availability-options {
    many-to-one {
      preferred-backup preferred-backup;
    }
  }
}
```

Load Balancing Network Address Translation Flows

Starting with Junos OS Release 11.4, Network Address Translation (NAT) has been programmed as a plug-in and is a function of load balancing and high availability. The plug-in runs on AMS infrastructure. All flows for translation are automatically distributed to different services PICs that are part of the AMS infrastructure. In case of failure of an active Multiservices PIC, the configured backup Multiservices PIC will take over the NAT pool resources of the failed PIC. The hashing method selected depends on the type of NAT. Using NAT on AMS infrastructure has a few limitations:

- NAT flows to failed PICs cannot be restored.
- There is no support for IPv6 flows.
- Twice NAT is not supported for load balancing.

See “[Example: Configuring Static Source Translation on AMS Infrastructure](#)” on page 616 for more details on configuring NAT flows for load balancing.

Related Documentation

- [Understanding Aggregated Multiservices Interfaces on page 599](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 608](#)
- [Example: Configuring Next-Hop Style Services on an Aggregated Multiservices Interface on page 613](#)

- [Example: Configuring Static Source Translation on AMS Infrastructure on page 616](#)

Example: Configuring an Aggregated Multiservices Interface (AMS)

- [Hardware and Software Requirements on page 608](#)
- [Overview on page 608](#)
- [Configuration on page 609](#)
- [Verification on page 612](#)

Hardware and Software Requirements

This example requires MX Series routers that have services interfaces installed in that and Junos OS Release 13.2 running on that.

Overview

The aggregated multiservices (AMS) interface configuration in Junos OS enables you to combine multiple services interfaces to create a bundle of interfaces that can function as a single interface. This example shows you how to configure an AMS interface, load-balancing options, member failure options, high availability settings on an AMS interface, and an interface-style service set configuration that uses the AMS interface.



NOTE: You cannot include MS-DPCs or other multiservices PICs in an AMS configuration that contains MS-MICs or MS-MPCs as member interfaces.

An MS-PIC contains only one interface, whereas the MS-MPC contains four interfaces. To utilize the entire MS-MPC in a single AMS bundle, all the four member interfaces need to be assigned to that AMS bundle.

Keep the following points in mind for every member interface (XLP chip) needs to be part of the AMS interface bundle:

- XLP-based line cards from the same MPC can be part of multiple AMS bundles.
- Multiple XLP chips from several MPCs can also be part of a single bundle (up to eight member interfaces in an AMS bundle, depending on the deployment requirement).
- It is not necessary that all the XLP chips from the same MS-MPC must be part of the same AMS bundle. Some of the XLP chips can be part of an AMS bundle, while other XLP chips can be standalone **ms-** interfaces or need not be configured. However, the same XLP chip cannot be part of two different AMS interfaces at the same time. For example, each XLP chip from the same MS-MPC can be grouped into four different AMS bundles, based on the deployment needs.
- A maximum of up to eight member interfaces can be assigned to an AMS bundle.

For more information about AMS interfaces, see [“Understanding Aggregated Multiservices Interfaces” on page 599](#).

Configuration

CLI Quick Configuration	To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.
Adding Member Interfaces	<pre>set interfaces ams0 load-balancing-options member-interface mams-0/0/0 set interfaces ams0 load-balancing-options member-interface mams-0/1/0 set interfaces ams0 load-balancing-options member-interface mams-1/0/0 set interfaces ams0 load-balancing-options member-interface mams-1/1/0 set interfaces ams0 load-balancing-options member-interface mams-2/0/0 set interfaces ams0 load-balancing-options member-interface mams-2/1/0</pre>
Configuring Logical Units	<pre>set interfaces ams0 unit 1 family inet</pre>
Configuring Member Failure Options	<pre>set interfaces ams0 load-balancing-options member-failure-options drop-member-traffic rejoin-timeout 300 set interfaces ams0 load-balancing-options member-failure-options drop-member-traffic enable-rejoin</pre>
Configuring High Availability Options	<pre>set interfaces ams0 load-balancing-options high-availability-options many-to-one preferred-backup mams-1/0/0</pre>
Configuring Service Set and Interface Services	<pre>set services service-set ams-ssl interface-service service-interface ams0.1 set services service-set ams-ssl interface-service load-balancing-options hash-keys ingress-key source-ip set services service-set ams-ssl interface-service load-balancing-options hash-keys egress-key destination-ip</pre>
Step-by-Step Procedure	<p>The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see <i>Using the CLI Editor in Configuration Mode</i> in the <i>CLI User Guide</i>.</p> <ol style="list-style-type: none"> 1. Create an aggregated multiservices interface and add member interfaces.



NOTE: You cannot configure the same mams to be part of two different AMS interfaces at the same time.

```
[edit]
user@router1# set interfaces ams0 load-balancing-options member-interface
mams-0/0/0
user@router1# set interfaces ams0 load-balancing-options member-interface
mams-0/1/0
user@router1# set interfaces ams0 load-balancing-options member-interface
mams-1/0/0
user@router1# set interfaces ams0 load-balancing-options member-interface
mams-1/1/0
user@router1# set interfaces ams0 load-balancing-options member-interface
mams-2/0/0
```

```
user@router1# set interfaces ams0 load-balancing-options member-interface  
mams-2/1/0
```

2. Configure logical units for the AMS interface.



NOTE: An AMS interface and its member interfaces cannot share the same logical interface units. For example, if one of the member interfaces has logical units 1 and 2 configured on it, you cannot configure logical units 1 and 2 for the AMS. Similarly, if you have configured logical units 3 and 4 on the AMS, you cannot configure those units on any of the member interfaces.

```
[edit interfaces]  
user@router1# set ams0 unit 1 family inet
```

3. Configure member failure options.

```
[edit interfaces ams0]  
user@router1# set load-balancing-options member-failure-options  
drop-member-traffic rejoin-timeout 300  
user@router1# set load-balancing-options member-failure-options  
drop-member-traffic enable-rejoin
```



NOTE: This example shows the `drop-member-traffic` configuration. However, if you would like to redistribute the traffic to other available members when one of the member interfaces goes down, you can include the `redistribute-all-traffic` statement instead of the `drop-member-traffic` statement.

The default behavior, when the `member-failure-options` configuration is not included, is to drop member traffic with a rejoin timeout of 120 seconds.

4. Configure the high-availability options.

```
[edit interfaces ams0]  
user@router1# set load-balancing-options high-availability-options many-to-one  
preferred-backup mams-1/0/0
```

5. Configure interface style services.

```
[edit services]  
user@router1# set service-set ams-ssl interface-service service-interface ams0.1  
user@router1# set service-set ams-ssl interface-service load-balancing-options  
hash-keys ingress-key source-ip  
user@router1# set service-set ams-ssl interface-service load-balancing-options  
hash-keys egress-key destination-ip
```

6. If you are done configuring the device, commit the configuration.

```
[edit]  
user@router1# commit
```

Table 25: Key Configuration Statements Used in this Example

Statement	Description
member-interface	Adds a member interface (mams) to the AMS bundle.
drop-member-traffic	Specifies that all traffic to a member be dropped in case the member interface fails.
rejoin-timeout	Specifies the time interval, in seconds, for the AMS to wait before declaring a member interface down. If the failed member comes back online during this period, it can rejoin the AMS and resume traffic forwarding. The range is 0 through 1000 seconds.
enable-rejoin	Specifies whether a failed interface be allowed to rejoin the AMS when it comes back online. If this statement is not included in the configuration, you must manually add the interface to the AMS when the interface is back online.
preferred-backup	Designates a member interface as the floating backup.
interface-services	Specifies a service interface, an AMS interface in this example, to handle interface services.
hash-keys	Specifies the load-balancing hash keys. You can configure the following hash key values: source-ip , destination-ip , iif (incoming interface), oif (outgoing interface), and protocol . NOTE: For services that require traffic symmetry, you must configure symmetrical hashing. Symmetrical hashing configuration ensures that both forward and reverse traffic are routed through the same member interface.

Results From the configuration mode, confirm your configuration by entering the **show interfaces ams0** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@router1# show interfaces ams0
load-balancing-options {
  member-interface mams-0/0/0;
  member-interface mams-0/1/0;
  member-interface mams-1/0/0;
  member-interface mams-1/1/0;
  member-interface mams-2/0/0;
  member-interface mams-2/1/0;
  member-failure-options {
    drop-member-traffic {
      rejoin-timeout 300;
      enable-rejoin;
    }
  }
}

```

```
high-availability-options {
  many-to-one {
    preferred-backup mams-1/0/0;
  }
}
unit 1 {
  family inet;
}

user@router1# show services
service-set ams-ssl {
  interface-service {
    service-interface ams0.1;
    load-balancing-options {
      hash-keys {
        ingress-key source-ip;
        egress-key destination-ip;
      }
    }
  }
}
```

Verification

Confirm that the configuration is working properly.

- [Verifying the AMS Configuration on page 612](#)

Verifying the AMS Configuration

Purpose Verify the AMS configuration and status of member interfaces.

Action From operational mode, enter the **show** command.

```
user@router1> show interfaces load-balancing detail
Load-balancing interfaces detail
Interface      : ams0
State          : Up
Last change    : 00:01:28
Member count   : 6
HA Model       : Many-to-One
Members       :
  Interface    Weight  State
  mams-0/0/0   10      Active
  mams-0/1/0   10      Active
  mams-1/0/0   10      Backup
  mams-1/1/0   10      Active
  mams-2/0/0   10      Active
  mams-2/1/0   10      Active
```

Meaning Shows that **ams0** has six member interfaces with a many-to-one backup configuration. Of the six member interfaces, five are in active state and one, **mams-1/0/0**, is in backup state.

- Related Documentation**
- [Example: Configuring Next-Hop Style Services on an Aggregated Multiservices Interface on page 613](#)
 - [Understanding Aggregated Multiservices Interfaces on page 599](#)

Example: Configuring Next-Hop Style Services on an Aggregated Multiservices Interface

- [Hardware and Software Requirements on page 613](#)
- [Overview on page 613](#)
- [Configuration on page 613](#)

Hardware and Software Requirements

MX Series routers with services interfaces installed and running Junos OS Release 13.2.

Overview

Starting with Release 13.2, Junos OS extends next-hop style services support to aggregated multiservices (AMS) interfaces. In releases earlier than 12.3, only interface style services configurations were supported on AMS interfaces.

The next-hop style services configuration on AMS interfaces is different from the interface style services configuration. For next-hop style services, the load-balancing hash keys are defined as part of the logical unit configuration of the AMS interface. For interface style services, the hash keys configuration falls under the service-set configuration.

This example explains the next-hop style services configuration on an AMS interface, and shows the verification steps to verify that the configuration is working correctly.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Configuring an aggregated multiservices interface

```
set interfaces ams0 load-balancing-options member-interface mams-1/0/0
set interfaces ams0 load-balancing-options member-interface mams-1/1/0
set interfaces ams0 load-balancing-options member-interface mams-2/0/0
set interfaces ams0 load-balancing-options member-interface mams-2/1/0
set interfaces ams0 unit 1 family inet
set interfaces ams0 unit 1 service-domain inside
set interfaces ams0 unit 2 family inet
set interfaces ams0 unit 2 service-domain outside
```

Configuring Routing Instances that Use AMS interfaces

```
set routing-instances ri-internal instance-type virtual-router
set routing-instances ri-internal interface ge-0/0/2.0
set routing-instances ri-internal interface ams0.1
set routing-instances ri-internal routing-options static route 22.22.22.0/24 next-hop ams0.1
set routing-instances ri-external instance-type virtual-router
set routing-instances ri-external interface ge-2/0/6.0
set routing-instances ri-external interface ams0.2
```

	<code>set routing-instances ri-external routing-options static route 0.0.0.0/0 next-hop ams0.2</code>
Configuring Hash Keys	<code>set interfaces ams0 unit 1 load-balancing-options hash-keys ingress-key source-ip protocol</code> <code>set interfaces ams0 unit 2 load-balancing-options hash-keys ingress-key destination-ip protocol</code>
Configure Next Hop Services	<code>set services service-set ams-test stateful-firewall-rules sfw1</code> <code>set services service-set ams-test next-hop-service inside-service-interface ams0.1</code> <code>set services service-set ams-test next-hop-service outside-service-interface ams0.2</code>

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see “*Using the CLI Editor in Configuration Mode*” in the *CLI User Guide*.

1. Configure an aggregated multiservices interface and the load-balancing options.

```
[edit interfaces ams0]
user@router1# set load-balancing-options member-interface mams-1/0/0
user@router1# set load-balancing-options member-interface mams-1/1/0
user@router1# set load-balancing-options member-interface mams-2/0/0
user@router1# set load-balancing-options member-interface mams-2/1/0
user@router1# set unit 1 family inet
user@router1# set unit 1 service-domain inside
user@router1# set unit 2 family inet
user@router1# set unit 2 service-domain outside
```

2. Configure routing instances that use the aggregated multiservices interfaces configured in the first step.

```
[edit routing-instances]
user@router1# set ri-internal instance-type virtual-router
user@router1# set ri-internal interface ge-0/0/2.0
user@router1# set ri-internal interface ams0.1
user@router1# set ri-internal routing-options static route 22.22.22.0/24 next-hop ams0.1
user@router1# set ri-external instance-type virtual-router
user@router1# set ri-external interface ge-2/0/6.0
user@router1# set ri-external interface ams0.2
user@router1# set ri-external routing-options static route 0.0.0.0/0 next-hop ams0.2
```

3. Configure hash keys for the aggregated multiservices interfaces.



NOTE: Unlike in the interface-style configuration where hash keys are defined in the service-set configuration, for next-hop services, the hash keys are specified in the AMS configuration under the logical units.

```
[edit interfaces ams0]
user@router1# set unit 1 load-balancing-options hash-keys ingress-key source-ip protocol
user@router1# set unit 2 load-balancing-options hash-keys ingress-key destination-ip protocol
```

4. Configure next-hop style services under the service-set configuration.

```
[edit services service-set ams-test]
```



```

user@router1# set stateful-firewall-rules sfw1
user@router1# set next-hop-service inside-service-interface ams0.1
user@router1# set next-hop-service outside-service-interface ams0.2

```

5. Commit the configuration.

```

[edit]
user@router1# commit

```

Results From the configuration mode, confirm your configuration by entering the **show interfaces ams0**, **show routing-instances**, and **show services service-set ams-test** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@router1# show interfaces ams0
load-balancing-options {
  member-interface mams-1/0/0;
  member-interface mams-1/1/0;
  member-interface mams-2/0/0;
  member-interface mams-2/1/0;
  member-failure-options {
    redistribute-all-traffic {
      enable-rejoin;
    }
  }
}
unit 1 {
  family inet;
  service-domain inside;
  load-balancing-options {
    hash-keys {
      ingress-key [ source-ip protocol ];
    }
  }
}
unit 2 {
  family inet;
  service-domain outside;
  load-balancing-options {
    hash-keys {
      ingress-key [ destination-ip protocol ];
    }
  }
}

user@router1# show routing-instances
ri-internal {
  instance-type virtual-router;
  interface ge-0/0/2.0;
  interface ams0.1
  routing-options {
    static {
      route 22.22.22.0/24 next-hop ams0.1;
    }
  }
}

```

```
ri-external {
  instance-type virtual-router;
  interface ge-2/0/6.0;
  interface ams0.2
  routing-options {
    static {
      route 0.0.0.0/0 next-hop ams0.2;
    }
  }
}

user@router1# show services service-set ams
stateful-firewall-rules sfw1;
next-hop-service {
  inside-service-interface ams0.1;
  outside-service-interface ams0.2;
}
```

- Related Documentation**
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 608](#)
 - [Understanding Aggregated Multiservices Interfaces on page 599](#)

Example: Configuring Static Source Translation on AMS Infrastructure

This example shows a static source translation configured on an AMS interface. The flows will be load balanced across member interfaces with this example.

Configure the AMS interface **ams0** with load balancing options.

```
[edit interfaces ams0]
load-balancing-options {
  member-interface mams-5/0/0;
  member-interface mams-5/1/0;
}
unit 1 {
  family inet;
}
unit 2 {
  family inet;
}
```

Configure hashing for the service set for both ingress and egress traffic.

```
[edit services service-set ss1]
interface-service {
  service-interface ams0.1;
  load-balancing-options {
    hash-keys {
      ingress-key destination-ip;
      egress-key source-ip;
    }
  }
}
```



NOTE: Hashing is determined based on whether the service set is applied on the ingress or egress interface.

Configure two NAT pools because you have configured two member interfaces for the AMS interface.

```
[edit services]
nat {
  pool p1 {
    address-range low 20.1.1.80 high 20.1.1.80;
  }
  pool p2 {
    address 20.1.1.81/32;
  }
}
```

Configure the NAT rule and translation.

```
[edit services]
nat {
  rule r1 {
    match-direction input;
    term t1 {
      from {
        source-address {
          20.1.1.2/32;
        }
      }
      then {
        translated {
          source-pool p1;
          translation-type {
            basic-nat44;
          }
        }
      }
    }
    term t1 {
      from {
        source-address {
          40.1.1.2/32;
        }
      }
      then {
        translated {
          source-pool p2;
          translation-type {
            basic-nat44;
          }
        }
      }
    }
  }
}
```



NOTE: A similar configuration can be applied for translation types `dynamic-nat44` and `napt-44`. Twice NAT cannot run on AMS infrastructure at this time.

**Related
Documentation**

- [Configuring Load Balancing on AMS Infrastructure on page 605](#)
- [Understanding Aggregated Multiservices Interfaces on page 599](#)

PART 11

Handling VoIP, HTTP, and Layer 2 Traffic

- [Handling VoIP Traffic Using Voice Services on page 621](#)
- [Handling HTTP Traffic Using HTTP Content Manager \(HCM\) on page 631](#)
- [Tunneling PPP Packets Across a Network Using Layer 2 Tunneling on page 637](#)

Handling VoIP Traffic Using Voice Services

- [Voice Services Overview on page 621](#)
- [Configuring Services Interfaces for Voice Services on page 622](#)
- [Configuring Encapsulation for Voice Services on page 625](#)
- [Configuring Network Interfaces for Voice Services on page 626](#)
- [Examples: Configuring Voice Services on page 627](#)

Voice Services Overview

Adaptive services interfaces include a voice services feature that allows you to specify interface type **lsq-fpc/pic/port** to accommodate voice over IP (VoIP) traffic. This interface uses compressed RTP (CRTP), which is defined in RFC 2508, *Compressing IP/UDP/RTP Headers for Low-Speed Serial Links*.

CRTP enables VoIP traffic to use low-speed links more effectively, by compressing the 40-byte IP/UDP/RTP header down to 2 to 4 bytes in most cases.

Voice services on the AS and MultiServices PICs support single-link PPP-encapsulated IPv4 traffic over the following physical interface types: ATM2, DS3, E1, E3, OC3, OC12, STM1, and T1, including the channelized versions of these interfaces.

Voice services do not require a separate service rules configuration.

Voice services also support LFI on Juniper Networks M Series Multiservice Edge routers, except the M320 router. For more information about configuring voice services, see [“Configuring Services Interfaces for Voice Services” on page 622](#).

For link services IQ interfaces (**lsq**) only, you can configure CRTP with multiclass MLPPP (MCML). MCML greatly simplifies packet ordering issues that occur when multiple links are used. Without MCML, all voice traffic belonging to a single flow is hashed to a single link in order to avoid packet ordering issues. With MCML, you can assign voice traffic to a high-priority class, and you can use multiple links. For more information about MCML support on link services IQ interfaces, see [“Configuring Link Services and CoS on Services PICs” on page 529](#).

Related Documentation

- [Configuring Services Interfaces for Voice Services on page 622](#)
- [Configuring Encapsulation for Voice Services on page 625](#)

- [Configuring Network Interfaces for Voice Services on page 626](#)
- [Examples: Configuring Voice Services on page 627](#)

Configuring Services Interfaces for Voice Services

You define voice service properties such as compression by configuring statements and values for a voice services interface, specified by the interface type **lsq**-. You can include the following statements:

```
encapsulation mlppp;
family inet {
    address address;
}
compression {
    rtp {
        f-max-period number;
        maximum-contexts number <force>;
        port {
            minimum port-number;
            maximum port-number;
        }
        queues [ queue-numbers ];
    }
}
fragment-threshold bytes;
```

You can include these statements at the following hierarchy levels:

- [edit interfaces (lsq | ls)-fpc/pic/port unit logical-unit-number]
- [edit logical-systems logical-system-name interfaces (lsq | ls)-fpc/pic/port unit logical-unit-number]

The following sections provide detailed instructions for configuring for voice services on services interfaces:

- [Configuring the Logical Interface Address for the MLPPP Bundle on page 622](#)
- [Configuring Compression of Voice Traffic on page 623](#)
- [Configuring Delay-Sensitive Packet Interleaving on page 624](#)
- [Example: Configuring Compression of Voice Traffic on page 624](#)

Configuring the Logical Interface Address for the MLPPP Bundle

To configure the logical address for the MLPPP bundle, include the **address** statement:

```
address address {
    ...
}
```

You can configure this statement at the following hierarchy levels:

- [edit interfaces (lsq | ls)-fpc/pic/port unit logical-unit-number family inet]

- [edit logical-systems *logical-system-name* interfaces (lsq | ls)-fpc/pic/port unit *logical-unit-number* family inet]

address specifies an IP address for the interface. AS and Multiservices PICs support only IP version 4 (IPv4) addresses, which are therefore configured under the **family inet** statement.

For information on other addressing properties you can configure that are not specific to service interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

Configuring Compression of Voice Traffic

You can specify how a services interface handles voice traffic compression by including the **compression** statement:

```
compression {
  rtp {
    f-max-period number;
    maximum-contexts number <force>;
    port {
      minimum port-number;
      maximum port-number;
    }
    queues [ queue-numbers ];
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces (lsq | ls)-fpc/pic/port unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces (lsq | ls)-fpc/pic/port unit *logical-unit-number*]

The following statements configure the indicated compression properties:

- **f-max-period *number***—Sets the maximum number of compressed packets to insert between the transmission of full headers. If you do not include the statement, the default is 255 packets.
- **maximum-contexts *number* <force>**—Specifies the maximum number of RTP contexts to accept during negotiation. The optional **force** statement requires the PIC to use the value specified for maximum RTP contexts, regardless of the negotiated value. This option enables interoperability with Junos OS Releases that base the RTP context value on link speed.
- **port, minimum *port-number*, and maximum *port-number***—Specify the lower and upper boundaries for a range of UDP destination port values on which RTP compression takes effect. Values for **port-number** can range from 0 through 65,535. RTP compression is applied to traffic transiting the ports within the specified range.
- **queues [*queue-numbers*]**—Specifies one or more of queues **q0**, **q1**, **q2**, and **q3**. RTP compression is applied to the traffic in the specified queues.



NOTE: If you specify both a port range and one or more queues, compression takes place if either condition is met.

Configuring Delay-Sensitive Packet Interleaving

When you configure CRTP, the software automatically enables link fragmentation and interleaving (LFI). LFI reduces excessive delays by fragmenting long packets into smaller packets and interleaving them with real-time frames. This allows real-time and non-real-time data frames to be carried together on lower-speed links without causing excessive delays to the real-time traffic. When the peer interface receives the smaller fragments, it reassembles the fragments into their original packet. For example, short delay-sensitive packets, such as packetized voice, can race ahead of larger delay-insensitive packets, such as common data packets.

By default, LFI is always active when you include the **compression rtp** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level. You control the operation of LFI indirectly by setting the **fragment-threshold** statement on the same logical interface. For example, if you include the **fragment-threshold 256** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level, all IP packets larger than 256 bytes are fragmented.

Example: Configuring Compression of Voice Traffic

Configure compression on a T1 interface with MLPPP encapsulation. Configure fragmentation for all IP packets larger than 128 bytes.

```
[edit interfaces]
t1-1/0/0 {
  unit 0 {
    family mlppp {
      bundle lsq-1/1/0.1;
    }
  }
}
lsq-1/1/0 {
  encapsulation mlppp;
  unit 1 {
    compression {
      rtp {
        port minimum 2000 maximum 64009;
      }
    }
    family inet {
      address 30.1.1.2/24;
    }
    fragment-threshold 128;
  }
}
```

Related Documentation

- [Voice Services Overview on page 621](#)

- [Configuring Encapsulation for Voice Services on page 625](#)
- [Configuring Network Interfaces for Voice Services on page 626](#)
- [Examples: Configuring Voice Services on page 627](#)

Configuring Encapsulation for Voice Services

Voice services interfaces support the following logical interface encapsulation types:

- Multilink Point-to-Point Protocol (MLPPP), which is the default encapsulation
- ATM2 IQ MLPPP over AAL5 LLC
- Frame Relay PPP

For general information on encapsulation, see the *Junos OS Network Interfaces Library for Routing Devices*. You can also configure physical interface encapsulation on voice services interfaces.

To configure voice services encapsulation, include the **encapsulation** statement:

encapsulation *type*;

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For voice services interfaces, the valid values for the **type** variable are **atm-mlppp-llc**, **frame-relay-ppp** or **multilink-ppp**.

You must also configure the physical interface with the corresponding encapsulation type, either Frame Relay or PPP. LSQ interfaces are supported by the following physical interface types: ATM2 IQ, DS3, E1, E3, OC3, OC12, STM1, and T1, including the channelized versions of these interfaces. For examples, see “[Examples: Configuring Voice Services](#)” on page 627.



NOTE: The only protocol type supported with **frame-relay-ppp** encapsulation is **family mlppp**.

Related Documentation

- [Voice Services Overview on page 621](#)
- [Configuring Services Interfaces for Voice Services on page 622](#)
- [Configuring Network Interfaces for Voice Services on page 626](#)
- [Examples: Configuring Voice Services on page 627](#)

Configuring Network Interfaces for Voice Services

To complete a voice services interface configuration, you need to configure the physical network interface with either MLPPP encapsulation and a voice services bundle or PPP encapsulation and a compression interface, as described in the following sections:

- [Configuring Voice Services Bundles with MLPPP Encapsulation on page 626](#)
- [Configuring the Compression Interface with PPP Encapsulation on page 626](#)

Configuring Voice Services Bundles with MLPPP Encapsulation

For voice services interfaces, you configure the link bundle as a channel. The physical interface is usually connected to networks capable of supporting MLPPP; the interface types supported for voice traffic are T1, E1, T3, E3, OC3, OC12, and STM1, including channelized versions of these interfaces.



NOTE:

For M Series routers and T Series routers, the following caveats apply:

- Maximum supported throughput on the bundle interfaces is 45 Mbps.
- Bundling of the logical interfaces under a T3 physical interface into the same or different bundles is not supported.

To configure a physical interface link for MLPPP, include the following statement:

```
bundle interface-name;
```

You can configure this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family mlppp]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family mlppp]**

When you configure **family mlppp**, no other protocol configuration is allowed. For more information on link bundles, see [“Configuring the Links in a Multilink or Link Services Bundle” on page 721](#).

Configuring the Compression Interface with PPP Encapsulation

To configure the physical interface for PPP encapsulation, you also need to specify the services interface to be used for voice compression: a Link Services IQ (**lsq-**) interface.

To configure the compression interface, include the **compression-device** statement:

```
compression-device interface-name;
```

You can configure this statement at the following hierarchy levels:

- **[edit interfaces (lsq | ls)-*fpc/pic/port* unit *logical-unit-number*]**

- `[edit logical-systems logical-system-name interfaces (lsq | ls)-fpc/pic/port unit logical-unit-number]`

**Related
Documentation**

- [Voice Services Overview on page 621](#)
- [Configuring Services Interfaces for Voice Services on page 622](#)
- [Configuring Encapsulation for Voice Services on page 625](#)
- [Examples: Configuring Voice Services on page 627](#)

Examples: Configuring Voice Services

Configure voice services using a T1 physical interface and MLPPP bundle encapsulation:

```
[edit interfaces]
t1-0/2/0:1 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle lsq-1/3/0.1;
    }
  }
}
lsq-1/3/0 {
  unit 1 {
    encapsulation mlppp;
    family inet {
      address 10.5.5.2/30;
    }
    compression {
      rtp {
        f-max-period 100;
        queues [ q1 q2 ];
        port {
          minimum 16384;
          maximum 32767;
        }
      }
    }
  }
  fragment-threshold 128;
}
```

Configure voice services using Frame Relay encapsulation without bundling:

```
[edit interfaces]
t1-1/0/0 {
  encapsulation frame-relay;
  unit 0 {
    dlci 100;
    encapsulation frame-relay-ppp;
    compression-device lsq-2/0/0.0;
  }
}
```

```
lsq-2/0/0 {  
  unit 0 {  
    compression {  
      rtp {  
        f-max-period 100;  
        queues [ q1 q2 ];  
        port {  
          minimum 16000;  
          maximum 32000;  
        }  
      }  
    }  
  }  
  family inet {  
    address 10.1.1.1/32;  
  }  
}
```

Configure voice services using an ATM2 physical interface (the corresponding class-of-service configuration is provided for illustration):

```
[edit interfaces]  
at-1/2/0 {  
  atm-options {  
    vpi 0;  
    pic-type atm2; # only ATM2 PICs are supported  
  }  
  unit 0 {  
    vci 0.69;  
    encapsulation atm-mlppp-llc;  
    family mlppp {  
      bundle lsq-1/3/0.10;  
    }  
  }  
  unit 1 {  
    vci 0.42;  
    encapsulation atm-mlppp-llc;  
    family mlppp {  
      bundle lsq-1/3/0.11;  
    }  
  }  
}  
lsq-1/3/0 {  
  unit 10 {  
    encapsulation multilink-ppp;  
  }  
  # Large packets need to be fragmented.  
  # Fragmentation can also be specified per forwarding class.  
  fragment-threshold 320;  
  compression {  
    rtp {  
      port minimum 2000 maximum 64009;  
    }  
  }  
}  
unit 11 {
```

```

    encapsulation multilink-ppp;
  }
  fragment-threshold 160;
  [edit class-of-service]
  scheduler-maps {
    sched {
      # Scheduling parameters apply to bundles on the AS or Multiservices PIC.
      # Unlike DS3/SONET interfaces, there is no need to create
      # a separate scheduler map for the ATM PIC. ATM defines
      # CoS constructs under the [edit interfaces at-fpc/pic/port] hierarchy.
      ...
    }
  }
  fragmentation-maps {
    fragmap {
      forwarding-class {
        ef {
          # In this example, voice is carried in the ef queue.
          # It is interleaved with bulk data.
          # Alternatively, you could use multiclass MLPPP to
          # carry multiple classes of traffic in different
          # multilink classes.
          no-fragmentation;
        }
      }
    }
  }
  interfaces {
    # Assign fragmentation and scheduling parameters to LSQ interfaces.
    lsq-1/3/0 {
      unit 0 {
        shaping-rate 512k;
        scheduler-map sched;
        fragmentation-map fragmap;
      }
      unit 1 {
        shaping-rate 128k;
        scheduler-map sched;
        fragmentation-map fragmap;
      }
    }
  }
}

```

**Related
Documentation**

- [Voice Services Overview on page 621](#)
- [Configuring Services Interfaces for Voice Services on page 622](#)
- [Configuring Encapsulation for Voice Services on page 625](#)
- [Configuring Network Interfaces for Voice Services on page 626](#)

Handling HTTP Traffic Using HTTP Content Manager (HCM)

- [HTTP Content Manager \(HCM\) on page 631](#)
- [HTTP URL Tracking and Policy Control for Client Requests on page 634](#)
- [Configuring HTTP URL Tracking and Policy Control on page 635](#)

HTTP Content Manager (HCM)

HTTP Content Management (HCM) is an application used for inspecting the HTTP traffic transmitted through port 80 (default) or any other port you use to transmit HTTP traffic. HCM can be installed on an MX-series router that is running the corresponding version of the Junos OS release. HCM inspects HTTP traffic even if the default port 80 is not used for HTTP traffic and is interoperable with ms, rms, and ams interface types. It supports fragmented HTTP request packets and GET, PUT, and POST requests.

Configuring the HTTP-Manager Package on the Router

1. Before you install the HTTP-Manager package on the router, ensure that you have the appropriate version of the HTTP-Manager package for the Junos OS image you are using on the router. When you have confirmed that you have the right package, use the request system software add command to install the HTTP-Manager package. You would need to restart the CLI after the package is installed.

```
user@router> request system software add http-manager-12.2R2-1-A1.2.tgz
```

```
NOTICE: Validating configuration against package-name.
```

```
NOTICE: Use the 'no-validate' option to skip this if desired.
```

```
Checking compatibility with configuration
```

```
Initializing...
```

```
WARNING: cli has been replaced by an updated version:
```

```
CLI release 12.2R2 built by builder on 2012-01-24 02:36:22 UTC
```

```
Restart cli using the new version ? [yes,no] (yes)
```

```
Restarting cli ...
```

2. When the CLI has restarted, use the **show version** command to see whether the HTTP-Manager packages are installed.

```
user@router> show version
```

```
...
```

```
HTTP-Manager Management Component [12.2R2-1-A1.2]
```

HTTP-Manager Dataplane Component [12.2R2-1-A1.2]

user@router>..

3. If you want to upgrade the Junos OS image on a router that has the HTTP-Manager package installed, you should first save and then delete the HTTP-Manager configuration from the router.
 - To view the HTTP-Manager configuration, use the **user@router>extension juniper-http-manager show <section>** command.
 - To delete the HTTP-Manager configuration from the router, use the **user@router>extension juniper-http-manager delete <section>** command.
 - Any remnant HTTP-Manager configuration left on the router will be deleted when the Junos OS image is upgraded. So, ensure that you have saved all necessary HTTP Content Management configurations.
 - To delete the HTTP-Manager package from the router, use the **user@router>>request system software delete <http-manager-package>** command.
 - Reinstall the HTTP-Manager package on the router after you upgrade the Junos OS image on the router.

root@aulavik> show version

Hostname: aulavik

Model: mx480

JUNOS Base OS boot [12.2R2]

JUNOS Base OS Software Suite [12.2R2]

JUNOS Kernel Software Suite [12.2R2]

JUNOS Crypto Software Suite [12.2R2]

JUNOS Packet Forwarding Engine Support (M/T Common) [12.2R2]

JUNOS Packet Forwarding Engine Support (MX Common) [12.2R2]

JUNOS Online Documentation [12.2R2]

JUNOS Voice Services Container package [12.2R2]

JUNOS Border Gateway Function package [12.2R2]

JUNOS Services AACL Container package [12.2R2]

JUNOS Services LL-PDF Container package [12.2R2]

JUNOS Services PTSP Container package [12.2R2]

JUNOS Services Stateful Firewall [12.2R2]

JUNOS Services NAT [12.2R2]

JUNOS Services Application Level Gateways [12.2R2]

JUNOS Services Captive Portal and Content Delivery Container package [12.2R2]

JUNOS Services RPM [12.2R2]

JUNOS Services HTTP Content Management package [12.2R2]

JUNOS Appld Services [12.2R2]

JUNOS IDP Services [12.2R2]

JUNOS Services Crypto [12.2R2]

JUNOS Services SSL [12.2R2]

JUNOS Services IPSec [12.2R2]

JUNOS Runtime Software Suite [12.2R2]

JUNOS Routing Software Suite [12.2R2]

HTTP-Manager Management Component [12.2R2-1-A1.2]

HTTP-Manager Dataplane Component [12.2R2-1-A1.2]

root@aulavik> configure

Entering configuration mode

[edit]

```

root@aulavik# extension juniper-http-manager show
## Last changed: 2012-06-07 13:21:36 PDT
services {
  http-manager {
    traceoptions {
      level all;
      flag all;
    }
  }
}

[edit]
root@aulavik# extension juniper-http-manager delete

[edit]
root@aulavik# extension juniper-http-manager show

[edit]
root@aulavik# commit
commit complete

[edit]
root@aulavik# exit
Exiting configuration mode

root@aulavik> request system software delete http-manager-services
Removing package 'http-manager-services' ...
Removing /opt/sdk/service-packages/http-manager-services ...
Removing http-manager-services-xlr-12.2R2-1-A1.2.tgz from /var/sw/pkg ...
Notifying mspd ...

root@aulavik> request system software delete http-manager-mgmt
Removing package 'http-manager-mgmt' ...
Reloading /config/juniper.conf.gz ...
Activating /config/juniper.conf.gz ...
mgd: commit complete
Restarting mgd ...
Restarting http-manager ...

WARNING: cli has been replaced by an updated version:
CLI release 11.4R3.7 built by builder on 2012-05-14 19:51:45 UTC
Restart cli using the new version ? [yes,no] (yes)

Restarting cli ...
root@aulavik>

root@aulavik> show version
Hostname: aulavik
Model: mx480
JUNOS Base OS boot [12.2R2]
JUNOS Base OS Software Suite [12.2R2]
JUNOS Kernel Software Suite [12.2R2]
JUNOS Crypto Software Suite [12.2R2]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.2R2]
JUNOS Packet Forwarding Engine Support (MX Common) [12.2R2]
JUNOS Online Documentation [12.2R2]
JUNOS Voice Services Container package [12.2R2]
JUNOS Border Gateway Function package [12.2R2]
JUNOS Services AACL Container package [12.2R2]
JUNOS Services LL-PDF Container package [12.2R2]

```

JUNOS Services PTSP Container package [12.2R2]
JUNOS Services Stateful Firewall [12.2R2]
JUNOS Services NAT [12.2R2]
JUNOS Services Application Level Gateways [12.2R2]
JUNOS Services Captive Portal and Content Delivery Container package [12.2R2]
JUNOS Services RPM [12.2R2]
JUNOS Services HTTP Content Management package [12.2R2]
JUNOS Appld Services [12.2R2]
JUNOS IDP Services [12.2R2]
JUNOS Services Crypto [12.2R2]
JUNOS Services SSL [12.2R2]
JUNOS Services IPSec [12.2R2]
JUNOS Runtime Software Suite [12.2R2]
JUNOS Routing Software Suite [12.2R2]

**Related
Documentation**

- [Configuring HTTP URL Tracking and Policy Control on page 635](#)
- [HTTP URL Tracking and Policy Control for Client Requests on page 634](#)
- *show services hcm statistics*

HTTP URL Tracking and Policy Control for Client Requests

The URL manipulation capability in the service plane allows an administrator to enter all the URLs associated with an action. An administrator can now enter URLs and actions in the service plane and apply them to the traffic associated with an interface and a subscriber by using the existing functions.



NOTE: The URL manipulation capability is supported only when the Junos OS Extension-Provider packages are installed and configured on the device.

In Junos OS releases earlier than 12.3, the extension-provider packages were variously referred to as Junos Services Framework (JSF), MP-SDK, and eJunos.

The HTTP URL manipulation allows routers to:

- Monitor HTTP transactions in the service plane and identify matching incoming HTTP requests with preconfigured URLs. When a matching HTTP request is found, the associated action and preconfigured URL are applied to the transaction.
- Maintain the statistics of HTTP requests that match the preconfigured URLs. Such statistics must be maintained per policy on a per-rule or per-term basis across all service sets.
- Display the current statistics for all visible HTTP-URL matches.
- Scale 50,000 HTTP transactions per second on a single MS-DPC network processing unit.

Guidelines for Configuring HTTP URL Monitoring for Client Requests

HTTP URL manipulation can be configured with service rules containing a sequence of terms. The URL rules are evaluated based on the longest prefix or suffix matches. The rules and terms are configured if:

- Multiple hostnames are specified in a match condition with an OR operator (that is, “any” is a potential match).
- Multiple request-URLs are specified in a match condition with an OR operator (that is, “any” is a potential match).
- The “discard” action causes HTTP requests matched to URLs to be dropped (stateful).
- The “accept” action causes HTTP requests matched to URLs to be allowed (stateful).
- The “log-request” action causes HTTP requests matched to URLs to be logged (fast log, once per transaction).
- The “count” action causes HTTP requests matched to URLs to be counted against the specified rule or term. The total number is the total of all service sets in which the rule or terms exists.
- The “discard” action can be combined only with the “log-request” action.
- The “accept” action can be combined with the “log-request” and “count” actions.
- If any given HTTP request matches more than one rule or term, the action applied is undeterministic and may be any of the matched rules or terms.
- The number of rules, terms, url-lists, and individual clauses in the rules, terms, and url-lists are limited by service-plane memory. No hard limits are imposed.
- If a hostname is not specified, then “*” (any) is assumed. Similarly, if a request-URL is not specified, then “*” (any) is assumed. Matching of hostnames and request-URLs to the HTTP request follows the same process described for url-lists. However, a match happens only when both the hostname and the request-URL match an entry in the same term of a rule.
- Multiple “url” and “url-list” clauses may be entered for the same “from {}” clause.

Configuring HTTP URL Tracking and Policy Control

To configure an HTTP URL Tracking and Policy Control, include the **url-rule *url-rule-name*** statement at the **[edit services hcm]** hierarchy level:

```
services {
  hcm {
    url-rule url-rule-name {
      term term-num {
        from {
          url-list url-list-name;
          url url_identifier {
            host hostname;
            request-url page-name;
          }
        }
      }
    }
  }
}
```

```
    }
    then {
        discard;
        accept;
        count;
        log-request;
    }
}
url-rule-set url-rule-set-name {
    url-rule rule1;
    url-rule rule2;
}
}
```

Each HCM rule consists of a set of terms, similar to a firewall filter. A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded.
- **then** statement—Specifies the actions and action modifiers to be performed by the router software.

**Related
Documentation**

- [HTTP Content Manager \(HCM\) on page 631](#)
- [HTTP URL Tracking and Policy Control for Client Requests on page 634](#)
- *show services hcm statistics*

Tunneling PPP Packets Across a Network Using Layer 2 Tunneling

- [Layer 2 Tunneling Protocol Overview on page 637](#)
- [L2TP Services Configuration Overview on page 638](#)
- [L2TP Minimum Configuration on page 639](#)
- [Configuring L2TP Tunnel Groups on page 641](#)
- [Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 646](#)
- [AS PIC Redundancy for L2TP Services on page 648](#)
- [Examples: Configuring L2TP Services on page 648](#)
- [Tracing L2TP Operations on page 652](#)

Layer 2 Tunneling Protocol Overview

L2TP is defined in RFC 2661, *Layer Two Tunneling Protocol (L2TP)*.

L2TP facilitates the tunneling of PPP packets across an intervening network in a way that is as transparent as possible to both end users and applications. It employs access profiles for group and individual user access, and uses authentication to establish secure connections between the two ends of each tunnel. Multilink PPP functionality is also supported.

The L2TP services are supported on the following routers only:

- M7i routers with AS PICs
- M10i routers with AS and MultiServices 100 PICs
- M120 routers with AS, MultiServices 100, and MultiServices 400 PICs
- On MX Series routers, the L2TP access concentrator (LAC) and L2TP network server (LNS) functions are supported only on MPCs; they are not supported on any services PIC or MS-DPC. For details about MPC support for L2TP, see the [MX Series Interface Module Reference](#)

For more information, see “[L2TP Services Configuration Overview](#)” on page 638.

- Related Documentation**
- [L2TP Services Configuration Overview on page 638](#)
 - [AS PIC Redundancy for L2TP Services on page 648](#)
 - [L2TP Minimum Configuration on page 639](#)
 - [Examples: Configuring L2TP Services on page 648](#)

L2TP Services Configuration Overview

The statements for configuring L2TP services are found at the following hierarchy levels:

- **[edit services l2tp tunnel-group *group-name*]**

The L2TP **tunnel-group** statement identifies an L2TP instance or L2TP server. Associated statements specify the local gateway address on which incoming tunnels and sessions are accepted, the Adaptive Services (AS) Physical Interface Card (PIC) that processes data for the sessions in this tunnel group, references to L2TP and PPP access profiles, and other attributes for configuring window sizes and timer values.

- **[edit interfaces *sp-fpc/pic/port* unit *logical-unit-number* dial-options]**

The **dial-options** statement includes configuration for the **l2tp-interface-id** statement and the **shared/dedicated** flag. The interface identifier associates a user session with a logical interface. Sessions can use either shared or dedicated logical interfaces. To run routing protocols, a session must use a dedicated logical interface.

- **[edit access profile *profile-name* client *name* l2tp]**

Tunnel profiles are defined at the **[edit access]** hierarchy level. Tunnel clients are defined with authentication, multilink negotiation and fragmentation, and other L2TP attributes in these profiles.

- **[edit access profile *profile-name* client *name* ppp]**

User profiles are defined at the **[edit access]** hierarchy level. User clients are defined with authentication and other PPP attributes in these profiles. These client profiles are used when local authentication is specified.

- **[edit access radius-server *address*]**

When you configure **authentication-order radius** at the **[edit access profile *profile-name*]** hierarchy level, you must configure a RADIUS service at the **[edit access radius-server]** hierarchy level.



NOTE: For information about L2TP LAC and LNS configurations on MX Series routers for subscriber access, see *L2TP for Subscriber Access Overview*.

- Related Documentation**
- [Layer 2 Tunneling Protocol Overview on page 637](#)
 - [AS PIC Redundancy for L2TP Services on page 648](#)
 - [L2TP Minimum Configuration on page 639](#)

- [Examples: Configuring L2TP Services on page 648](#)

L2TP Minimum Configuration

To configure L2TP services, you must perform at least the following tasks:

- Define a tunnel group at the **[edit services l2tp]** hierarchy level with the following attributes:
 - **l2tp-access-profile**—Profile name for the L2TP tunnel.
 - **ppp-access-profile**—Profile name for the L2TP user.
 - **local-gateway**—Address for the L2TP tunnel.
 - **service-interface**—AS PIC interface for the L2TP service.
 - Optionally, you can configure **traceoptions** for debugging purposes.

The following example shows a minimum configuration for a tunnel group with trace options:

```
[edit services l2tp]
tunnel-group finance-lns-server {
  l2tp-access-profile westcoast_bldg_1_tunnel;
  ppp-access-profile westcoast_bldg_1;
  local-gateway {
    address 10.21.255.129;
  }
  service-interface sp-1/3/0;
}
traceoptions {
  flag all;
  filter {
    protocol udp;
    protocol l2tp;
    protocol ppp;
    protocol radius;
  }
}
```

- At the **[edit interfaces]** hierarchy level:
 - Identify the physical interface at which L2TP tunnel packets enter the router, for example **ge-0/3/0**.
 - Configure the AS PIC interface with **unit 0 family inet** defined for IP service, and configure another logical interface with **family inet** and the **dial-options** statement.

The following example shows a minimum interfaces configuration for L2TP:

```
[edit interfaces]
ge-0/3/0 {
  unit 0 {
    family inet {
      address 10.58.255.129/28;
    }
  }
}
```

```

    }
  }
}
sp-1/3/0 {
  unit 0 {
    family inet;
  }
  unit 20 {
    dial-options {
      l2tp-interface-id test;
      shared;
    }
    family inet;
  }
}

```

- At the **[edit access]** hierarchy level:
 - Configure a tunnel profile. Each client specifies a unique L2TP Access Concentrator (LAC) name with an **interface-id** value that matches the one configured on the AS PIC interface unit; **shared-secret** is authentication between the LAC and the L2TP Network Server (LNS).
 - Configure a user profile. If RADIUS is used as the authentication method, it needs to be defined.
 - Define the RADIUS server with an IP address, port, and authentication data shared between the router and the RADIUS server.



NOTE: When the L2TP Network Server (LNS) is configured with RADIUS authentication, the default behavior is to accept the preferred RADIUS-assigned IP address. Previously, the default behavior was to accept and install the nonzero peer IP address that came into the IP-Address option of the IPCP Configuration Request packet.

- Optionally, you can define a group profile for common attributes, for example **keepalive 0** to turn off keepalive messages.

The following example shows a minimum profiles configuration for L2TP:

```

[edit access]
group-profile westcoast_users {
  ppp {
    keepalive 0;
  }
}
profile westcoast_bldg_1_tunnel {
  client production {
    l2tp {
      interface-id test;
      shared-secret "$9$n8HX6A01RhVl1R"; # SECRET-DATA
    }
    user-group-profile westcoast_users;
  }
}

```

```

}
profile westcoast_bldg_1 {
  authentication-order radius;
}
radius-server {
  192.168.65.63 {
    port 1812;
    secret "$9$Vyb4ZHkPQ39mf9pORlexNdbgoZUjqP5"; # SECRET-DATA
  }
}

```

- Related Documentation**
- [L2TP Services Configuration Overview on page 638](#)
 - [Configuring L2TP Tunnel Groups on page 641](#)
 - [Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 646](#)
 - [Tracing L2TP Operations on page 652](#)
 - [Examples: Configuring L2TP Services on page 648](#)

Configuring L2TP Tunnel Groups

To establish L2TP service on a router, you need to identify an L2TP tunnel group and specify a number of values that define which access profiles, interface addresses, and other properties to use in creating a tunnel. To identify the tunnel group, include the **tunnel-group** statement at the **[edit services l2tp]** hierarchy level:

```

tunnel-group group-name {
  hello-interval seconds;
  hide-avps;
  l2tp-access-profile profile-name;
  local-gateway address {
    address address;
    gateway-name gateway-name;
  }
  maximum-send-window packets;
  ppp-access-profile profile-name;
  receive-window packets;
  retransmit-interval seconds;
  service-interface interface-name;
  syslog {
    host hostname {
      services severity-level;
      facility-override facility-name;
      log-prefix prefix-value;
    }
  }
  tunnel-timeout seconds;
}

```



NOTE: If you delete a tunnel group or mark it inactive, all L2TP sessions in that tunnel group are terminated. If you change the value of the `local-gateway` address or the `service-interface` statement, all L2TP sessions using those settings are terminated. If you change or delete other statements at the `[edit services l2tp tunnel-group group-name]` hierarchy level, new tunnels you establish will use the updated values but existing tunnels and sessions are not affected.

The following sections explain how to configure L2TP tunnel groups:

- [Configuring Access Profiles for L2TP Tunnel Groups on page 642](#)
- [Configuring the Local Gateway Address and PIC on page 642](#)
- [Configuring Window Size for L2TP Tunnels on page 643](#)
- [Configuring Timers for L2TP Tunnels on page 643](#)
- [Hiding Attribute-Value Pairs for L2TP Tunnels on page 644](#)
- [Configuring System Logging of L2TP Tunnel Activity on page 644](#)

Configuring Access Profiles for L2TP Tunnel Groups

To validate L2TP connections and session requests, you set up access profiles by configuring the `profile` statement at the `[edit access]` hierarchy level. You need to configure two types of profiles:

- L2TP tunnel access profile, which validates all L2TP connection requests to the specified local gateway address
- PPP access profile, which validates all PPP session requests through L2TP tunnels established to the local gateway address

For more information on configuring the profiles, see the *Junos OS Administration Library for Routing Devices*. A profile example is included in “[Examples: Configuring L2TP Services](#)” on page 648.

To associate the profiles with a tunnel group, include the `l2tp-access-profile` and `ppp-access-profile` statements at the `[edit services l2tp tunnel-group group-name]` hierarchy level:

```
l2tp-access-profile profile-name;  
ppp-access-profile profile-name;
```

Configuring the Local Gateway Address and PIC

When you configure an L2TP group, you must also define a local address for the L2TP tunnel connections and the AS PIC that processes the requests:

- To configure the local gateway IP address, include the `address` statement at the `[edit services l2tp tunnel-group group-name local-gateway]` hierarchy level:

```
address address;
```

- To configure the AS PIC, include the **service-interface** statement at the **[edit services l2tp tunnel-group group-name]** hierarchy level:

```
service-interface sp-fpc/pic/port;
```

You can optionally specify the logical unit number along with the service interface. If specified, the unit is used as a logical interface representing PPP sessions negotiated using this profile.



NOTE: If you change the local gateway address or the service interface configuration, all L2TP sessions using those settings are terminated.

Dynamic class-of-service (CoS) functionality is supported on L2TP LNS sessions or L2TP sessions with ATM VCs, as long as the L2TP session is configured to use an IQ2 PIC on the egress interface. For more information, see the *Class of Service Feature Guide for Routing Devices*.

Configuring Window Size for L2TP Tunnels

You can configure the maximum window size for packet processing at each end of the L2TP tunnel:

- The receive window size limits the number of concurrent packets the server processes. By default, the maximum is 16 packets. To change the window size, include the **receive-window** statement at the **[edit services l2tp tunnel-group group-name]** hierarchy level:

```
receive-window packets;
```

- The maximum-send window size limits the other end's receive window size. The information is transmitted in the receive window size attribute-value pair. By default, the maximum is 32 packets. To change the window size, include the **maximum-send-window** statement at the **[edit services l2tp tunnel-group group-name]** hierarchy level:

```
maximum-send-window packets;
```

Configuring Timers for L2TP Tunnels

You can configure the following timer values that regulate L2TP tunnel processing:

- Hello interval—If the server does not receive any messages within a specified time interval, the router software sends a hello message to the tunnel's remote peer. By default, the interval length is 60 seconds. If you configure a value of 0, no hello messages are sent. To configure a different value, include the **hello-interval** statement at the **[edit services l2tp tunnel-group group-name]** hierarchy level:

```
hello-interval seconds;
```

- Retransmit interval—By default, the retransmit interval length is 30 seconds. To configure a different value, include the **retransmit-interval** statement at the **[edit services l2tp tunnel-group group-name]** hierarchy level:

```
retransmit-interval seconds;
```

- Tunnel timeout—If the server cannot send any data through the tunnel within a specified time interval, it assumes that the connection with the remote peer has been lost and deletes the tunnel. By default, the interval length is 120 seconds. To configure a different value, include the **tunnel-timeout** statement at the **[edit services l2tp tunnel-group group-name]** hierarchy level:

```
tunnel-timeout seconds;
```

Hiding Attribute-Value Pairs for L2TP Tunnels

Once an L2TP tunnel has been established and the connection authenticated, information is encoded by means of attribute-value pairs. By default, this information is not hidden. To hide the attribute-value pairs once the shared secret is known, include the **hide-avps** statement at the **[edit services l2tp tunnel-group group-name]** hierarchy level:

```
hide-avps;
```

Configuring System Logging of L2TP Tunnel Activity

You can specify properties that control how system log messages are generated for L2TP services.

To configure interface-wide default system logging values, include the **syslog** statement at the **[edit services l2tp tunnel-group group-name]** hierarchy level:

```
syslog {
  host hostname {
    services severity-level;
    facility-override facility-name;
    log-prefix prefix-value;
  }
}
```

Configure the **host** statement with a hostname or IP address that specifies the system log target server. The hostname **local** directs system log messages to the Routing Engine. For external system log servers, the hostname must be reachable from the same routing instance to which the initial data packet (that triggered session establishment) is delivered. You can specify only one system logging hostname.

Table 26 on page 644 lists the severity levels that you can specify in configuration statements at the **[edit services l2tp tunnel-group group-name syslog host hostname]** hierarchy level. The levels from **emergency** through **info** are in order from highest severity (greatest effect on functioning) to lowest.

Table 26: System Log Message Severity Levels

Severity Level	Description
any	Includes all severity levels
emergency	System panic or other condition that causes the router to stop functioning

Table 26: System Log Message Severity Levels (*continued*)

Severity Level	Description
alert	Conditions that require immediate correction, such as a corrupted system database
critical	Critical conditions, such as hard drive errors
error	Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels
warning	Conditions that warrant monitoring
notice	Conditions that are not errors but might warrant special handling
info	Events or nonerror conditions of interest

We recommend setting the system logging severity level to **error** during normal operation. To monitor PIC resource usage, set the level to **warning**. To gather information about an intrusion attack when an intrusion detection system error is detected, set the level to **notice** for a specific service set. To debug a configuration or log Network Address Translation (NAT) events, set the level to **info**.

For more information about system log messages, see the *Junos OS System Log Messages Reference*.

To use one particular facility code for all logging to the specified system log host, include the **facility-override** statement at the **[edit services l2tp tunnel-group group-name syslog host hostname]** hierarchy level:

```
facility-override facility-name;
```

The supported facilities include: **authorization**, **daemon**, **ftp**, **kernel**, **user**, and **local0** through **local7**.

To specify a text prefix for all logging to this system log host, include the **log-prefix** statement at the **[edit services l2tp tunnel-group group-name syslog host hostname]** hierarchy level:

```
log-prefix prefix-text;
```

Related Documentation

- [L2TP Services Configuration Overview on page 638](#)
- [L2TP Minimum Configuration on page 639](#)
- [Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 646](#)
- [Tracing L2TP Operations on page 652](#)
- [Examples: Configuring L2TP Services on page 648](#)

Configuring the Identifier for Logical Interfaces that Provide L2TP Services

You can configure L2TP services on adaptive services interfaces on M7i, M10i, M120, and MX Series routers only. You must configure the logical interface to be dedicated or shared. If a logical interface is dedicated, it can represent only one session at a time. A shared logical interface can have multiple sessions.

To configure the logical interface, include the **l2tp-interface-id** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* *dial-options*]** hierarchy level:

```
l2tp-interface-id name;  
(dedicated | shared);
```

The **l2tp-interface-id** name configured on the logical interface must be replicated at the **[edit access profile *name*]** hierarchy level:

- For a user-specific identifier, include the **l2tp-interface-id** statement at the **[edit access profile *name* ppp]** hierarchy level.
- For a group identifier, include the **l2tp-interface-id** statement at the **[edit access profile *name* l2tp]** hierarchy level.

You can configure multiple logical interfaces with the same interface identifier, to be used as a pool for several users. For more information on configuring access profiles, see the *Junos OS Administration Library for Routing Devices*.



NOTE: If you delete the **dial-options** statement settings configured on a logical interface, all L2TP sessions running on that interface are terminated.

Example: Configuring Multilink PPP on a Shared Logical Interface

Multilink PPP is supported on either shared or dedicated logical interfaces. The following example can be used to configure many multilink bundles on a single shared interface:

```
interfaces {  
  sp-1/3/0 {  
    traceoptions {  
      flag all;  
    }  
    unit 0 {  
      family inet;  
    }  
    unit 20 {  
      dial-options {  
        l2tp-interface-id test;  
        shared;  
      }  
      family inet;  
    }  
  }  
}
```



```

access {
  profile t {
    client test {
      l2tp {
        interface-id test;
        multilink;
        shared-secret "$9$n8HX6A01RhLvL1R"; # SECRET-DATA
      }
    }
  }
  profile u {
    authentication-order radius;
  }
  radius-server {
    192.168.65.63 {
      port 1812;
      secret "$9$Vyb4ZHkPQ39mf9pORlexNdbgoZUjqP5"; # SECRET-DATA
    }
  }
}
services {
  l2tp {
    tunnel-group 1 {
      l2tp-access-profile t;
      ppp-access-profile u;
      local-gateway {
        address 10.70.1.1;
      }
      service-interface sp-1/3/0;
    }
    traceoptions {
      flag all;
      debug-level packet-dump;
      filter {
        protocol l2tp;
        protocol ppp;
        protocol radius;
      }
    }
  }
}

```

**Related
Documentation**

- [L2TP Services Configuration Overview on page 638](#)
- [L2TP Minimum Configuration on page 639](#)
- [Configuring L2TP Tunnel Groups on page 641](#)
- [Tracing L2TP Operations on page 652](#)
- [Examples: Configuring L2TP Services on page 648](#)

AS PIC Redundancy for L2TP Services

L2TP services support AS PIC redundancy. To configure redundancy, you specify a redundancy services PIC (**rsp**) interface in which the primary AS PIC is active and a secondary AS PIC is on standby. If the primary AS PIC fails, the secondary PIC becomes active, and all service processing is transferred to it. If the primary AS PIC is restored, it remains in standby and does not preempt the secondary AS PIC; you need to manually restore the services to the primary PIC. To determine which PIC is currently active, issue the **show interfaces redundancy** command.



NOTE: On L2TP, the only service option supported is *warm standby*, in which one backup PIC supports multiple working PICs. Recovery times are not guaranteed, because the configuration must be completely restored on the backup PIC after a failure is detected. The tunnels and sessions are torn down upon switchover and need to be restarted by the LAC and PPP client, respectively. However, configuration is preserved and available on the new active PIC, although the protocol state needs to be reestablished.

As with the other AS PIC services that support warm standby, you can issue the **request interfaces (revert | switchover)** command to manually switch between primary and secondary L2TP interfaces.

For more information, see [“Configuring AS or Multiservices PIC Redundancy” on page 41](#). For an example configuration, see [“Examples: Configuring L2TP Services” on page 648](#). For information on operational mode commands, see the [CLI Explorer](#).

Related Documentation

- [Layer 2 Tunneling Protocol Overview on page 637](#)
- [L2TP Services Configuration Overview on page 638](#)
- [Configuring AS or Multiservices PIC Redundancy on page 41](#)
- [L2TP Minimum Configuration on page 639](#)
- [Examples: Configuring L2TP Services on page 648](#)

Examples: Configuring L2TP Services

Configure L2TP with multiple group and user profiles and a pool of logical interfaces for concurrent tunnel sessions:

```
[edit access]
address-pool customer_a {
  address 10.1.1.1/32;
}
address-pool customer_b {
  address-range low 10.2.2.1 high 10.2.3.2;
}
group-profile sunnyvale_users {
  ppp {
```

```

        framed-pool customer_a;
        idle-timeout 15;
        primary-dns 192.168.65.1;
        secondary-dns 192.168.65.2;
        primary-wins 192.168.65.3;
        secondary-wins 192.168.65.4;
        interface-id west;
    }
}
group-profile eastcoast_users {
    ppp {
        framed-pool customer_b;
        idle-timeout 20;
        primary-dns 192.168.65.5;
        secondary-dns 192.168.65.6;
        primary-wins 192.168.65.7;
        secondary-wins 192.168.65.8;
        interface-id east;
    }
}
group-profile sunnyvale_tunnel {
    l2tp {
        maximum-sessions-per-tunnel 100;
        interface-id west_shared;
    }
}
group-profile east_tunnel {
    l2tp {
        maximum-sessions-per-tunnel 125;
        interface-id east_shared;
    }
}
profile sunnyvale_bldg_1 {
    client white {
        chap-secret "$9$3s2690leK8X7VKM7VwgaJn/Ctu1hclv87Ct87"; # SECRET-DATA
        ppp {
            idle-timeout 22;
            primary-dns 192.168.65.1;
            framed-ip-address 10.12.12.12/32;
            interface-id east;
        }
        group-profile sunnyvale_users;
    }
    client blue {
        chap-secret "$9$eq1KWxbwgZUHNdjmqmTF3uO1Rhr-dsoJDNd"; # SECRET-DATA
        group-profile sunnyvale_users;
    }
    authentication-order password;
}
profile sunnyvale_bldg_1_tunnel {
    client test {
        l2tp {
            shared-secret "$9$r3HKvLg4ZUDkX7JGjif5pOBIRS8LN"; # SECRET-DATA
            maximum-sessions-per-tunnel 75;
            interface-id west_shared;
            ppp-authentication chap;
        }
    }
}

```

```
    }
    group-profile sunnyvale_tunnel;
}
client production {
    l2tp {
        shared-secret
            "$9$R2QErV8X-goGylVwg4jiTz36/t0BEleWFnRhrlXxbs2aJDHqf3nCP5";
        ppp-authentication chap;
    }
    group-profile sunnyvale_tunnel;
}
}
[edit services]
l2tp {
    tunnel-group finance-lns-server {
        l2tp-access-profile sunnyvale_bldg_1_tunnel;
        ppp-access-profile sunnyvale_bldg_1;
        local-gateway {
            address 10.1.117.3;
        }
        service-interface sp-1/3/0;
        receive-window 1500;
        maximum-send-window 1200;
        retransmit-interval 5;
        hello-interval 15;
        tunnel-timeout 55;
    }
    traceoptions {
        flag all;
    }
}
[edit interfaces sp-1/3/0]
unit 0 {
    family inet;
}
unit 10 {
    dial-options {
        l2tp-interface-id foo-user;
        dedicated;
    }
    family inet;
}
unit 11 {
    dial-options {
        l2tp-interface-id east;
        dedicated;
    }
    family inet;
}
unit 12 {
    dial-options {
        l2tp-interface-id east;
        dedicated;
    }
    family inet;
}
```

```

unit 21 {
    dial-options {
        l2tp-interface-id west;
        dedicated;
    }
    family inet;
}
unit 30 {
    dial-options {
        l2tp-interface-id west_shared;
        shared;
    }
    family inet;
}
unit 40 {
    dial-options {
        l2tp-interface-id east_shared;
        shared;
    }
    family inet;
}

```

Configure L2TP redundancy:

```

interfaces {
    rsp0 {
        redundancy-options {
            primary sp-0/0/0;
            secondary sp-1/3/0;
        }
        unit 0 {
            family inet;
        }
        unit 11 {
            dial-options {
                l2tp-interface-id east_shared;
                shared;
            }
            family inet;
        }
    }
}

```

Related Documentation

- [L2TP Services Configuration Overview on page 638](#)
- [L2TP Minimum Configuration on page 639](#)
- [Configuring L2TP Tunnel Groups on page 641](#)
- [Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 646](#)
- [Tracing L2TP Operations on page 652](#)

Tracing L2TP Operations

Tracing operations track all AS PIC operations and record them in a log file in the `/var/log` directory. By default, this file is named `/var/log/l2tpd`.



NOTE: This topic refers to tracing L2TP LNS operations on M Series routers. To trace L2TP LAC operations on MX Series routers, see *Tracing L2TP Operations for Subscriber Access*.

To trace L2TP operations, include the **traceoptions** statement at the **[edit services l2tp]** hierarchy level:

```
traceoptions {
  debug-level level;
  file <filename> <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  filter {
    protocol name;
    user-name username;
  }
  flag flag;
  interfaces interface-name {
    debug-level severity;
    flag flag;
  }
  level (all | error | info | notice | verbose | warning);
  no-remote-trace;
}
```

You can specify the following L2TP tracing flags:

- **all**—Trace everything.
- **configuration**—Trace configuration events.
- **protocol**—Trace routing protocol events.
- **routing-socket**—Trace routing socket events.
- **rpd**—Trace routing protocol process events.

You can specify a trace level for PPP, L2TP, RADIUS, and User Datagram Protocol (UDP) tracing. To configure a trace level, include the **debug-level** statement at the **[edit services l2tp traceoptions]** hierarchy level and specify one of the following values:

- **detail**—Detailed debug information
- **error**—Errors only
- **packet-dump**—Packet decoding information

You can filter by protocol. To configure filters, include the **filter protocol** statement at the **[edit services l2tp traceoptions]** hierarchy level and specify one or more of the following protocol values:

- **ppp**
- **l2tp**
- **radius**
- **udp**

To implement filtering by protocol name, you must also configure either **flag protocol** or **flag all**.

You can also configure traceoptions for L2TP on a specific adaptive services interface. To configure per-interface tracing, include the **interfaces** statement at the **[edit services l2tp traceoptions]** hierarchy level:

```
interfaces interface-name {
  debug-level level;
  flag flag;
}
```



NOTE: Implementing traceoptions consumes CPU resources and affects the packet processing performance.

You can specify the **debug-level** and **flag** statements for the interface, but the options are slightly different from the general L2TP traceoptions. You specify the debug level as **detail**, **error**, or **extensive**, which provides complete PIC debug information. The following flags are available:

- **all**—Trace everything.
- **ipc**—Trace L2TP Inter-Process Communication (IPC) messages between the PIC and the Routing Engine.
- **packet-dump**—Dump each packet's content based on debug level.
- **protocol**—Trace L2TP, PPP, and multilink handling.
- **system**—Trace packet processing on the PIC.

Related Documentation

- [L2TP Services Configuration Overview on page 638](#)
- [L2TP Minimum Configuration on page 639](#)
- [Configuring L2TP Tunnel Groups on page 641](#)
- [Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 646](#)
- [Examples: Configuring L2TP Services on page 648](#)

PART 12

Configuring Application Aware Services Interfaces

- [Configuring Stateless, Rule-Based Services Using Application-Aware Access Lists on page 657](#)
- [Grouping Applications Together Using APPID on page 669](#)
- [Detecting Suspicious and Anomalous Network Traffic Using IDP on page 691](#)
- [Collecting Statistics and Tracking Data Using L-PDF on page 697](#)

Configuring Stateless, Rule-Based Services Using Application-Aware Access Lists

- [AACL Overview on page 657](#)
- [Best-Effort Application Identification of DPI-Serviced Flows on page 658](#)
- [Configuring AACL Rules on page 661](#)
- [Example: Configuring AACL Rules on page 666](#)
- [Configuring AACL Rule Sets on page 666](#)
- [Configuring Logging of AACL Flows on page 667](#)

AACL Overview



NOTE: Starting with Junos OS Release 12.1, all interface-style services are supported for dynamic Point-to-Point Protocol over Ethernet (PPPoE) subscribers on all MX Series routers with modular Modular Port Concentrators (MPCs).

The application-aware access list (AACL) service adds support for a new service that uses application names and groups as matching criteria for filtering traffic. AACL is a stateless, rules-based service that must be combined with application identification to enable policies to be applied to flows based on application and application group membership in addition to traditional packet matching rules. It is supported on MX Series routers equipped with Multiservices DPCs and on M120 or M320 routers equipped with Multiservices 400 PICs. Starting with Junos OS Release 11.3, AACL is supported on T320, T640, and T1600 routers also.

AACL is configured in a similar way to other rules-based services such as Network Address Translation (NAT), class of service (CoS), and stateful firewall. To configure AACL, include rule specifications for match criteria and actions at the **[edit services aacl]** hierarchy level. You can chain AACL rules along with other service rules by including them in a service-set definition at the **[edit services service-set]** hierarchy level, as previously documented.

There is one pair of related operational commands, **show/clear application-aware-access-list statistics**.

For more information on the CLI configuration, see the *Application Aware Services Interfaces Feature Guide for Routing Devices*. For more information on the operational command, see the [CLI Explorer](#).



NOTE: Because the Junos OS extension-provider package framework lacks aggressive constraint checks, you should not set the `policy-db-size` statement at the `[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]` hierarchy level to a high value. For Junos Application Aware (previously known as dynamic application awareness) configurations, the recommended values for the extension-provider options at this hierarchy level are as follows:

- `control-cores = 1`
- `data-cores = 7`
- `object-cache-size = 1280` (for Multiservices 400 PIC and Multiservices DPC)
- `policy-db-size = 200`
- Include these package values: `jservices-idp`, `jservices-appid`, `jservices-llpdf`, `jservices-aacl`

For more information about this configuration, see the following topics in the *SDK Applications Configuration Guide and Command Reference*:

- *Configuring Control and Data Cores*
- *Configuring Memory Settings*
- *Configuring Packages on the PIC*

**Related
Documentation**

- [Configuring AACL Rules on page 661](#)
- [Configuring AACL Rule Sets on page 666](#)
- [Configuring Logging of AACL Flows on page 667](#)
- [Example: Configuring AACL Rules on page 666](#)

Best-Effort Application Identification of DPI-Serviced Flows

This topic describes the following information:

- [Features that Support Application-Level Filtering on page 659](#)
- [Best-Effort Application Determination on page 659](#)
- [APPID, AACL, and L-PDF Processing in Preconvergence Scenarios on page 659](#)

Features that Support Application-Level Filtering

On MX Series routers equipped with Multiservices DPCs and M120 or M320 routers equipped with Multiservices 400 PICs, Intrusion Detection and Prevention (IDP) is accomplished by Deep Packet Inspection (DPI) of TCP, UDP, and ICMP flows. The application identification (APPID) feature defines applications as members of application groups in TCP/UDP/ICMP traffic. IDP depends on APPID for identification and detection of some Layer 7 applications.

The application-aware access list (AACL) service uses application names and groups as matching criteria for filtering traffic. The service defines the applications and application groups for which statistics are collected for a specific user or interface.

The local policy decision function (L-PDF) enables you to configure properties for statistics output. L-PDF supports a process that regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

Best-Effort Application Determination

Typically, APPID conclusively determines the Layer 7 application associated with a given DPI-serviced flow. In these cases, the application identification is final. Occasionally, APPID is only able to make an initial, inconclusive determination of the Layer 7 application associated with a given flow. This is referred to as a "best-effort" application identification. In such cases, the APPID process continues processing packets on that flow and might subsequently make a conclusive determination of the application associated with that flow. In some cases of best-effort application identification, the flow ends before a final application determination can be made.

APPID, AACL, and L-PDF Processing in Preconvergence Scenarios

The following sections describe APPID, AACL, and L-PDF processing in various stages of application identification for a DPI-serviced flow of TCP/UDP/ICMP traffic.

- [Prior to a Final or Best-Effort Application Identification on page 659](#)
- [Upon Best-Effort Application Identification on page 660](#)
- [While Application Identification Is on a Best-Effort Basis on page 660](#)
- [If a Flow Ends Before an Application Identification Is Made on page 660](#)
- [If a Flow Ends While Application Identification on a Best-Effort Basis on page 660](#)

Prior to a Final or Best-Effort Application Identification

During the time that APPID has not yet made either a final or best-effort determination of the application associated with a given flow, the flow does not contribute to any per-subscriber or per-application statistics collection.

The output of the following operational mode commands includes flows for which APPID has not yet made either a final or best-effort determination of the associated application:

- **show services local-policy-decision-function flows** (interface *interface-name* | subscriber *subscriber-name*)
- **show services application-aware-access-list flows** (interface *interface-name* | subscriber *subscriber-name*)

In the command output, the **Action** field displays "accept" and the **Application** or **Application group** field displays "unknown" for a flow for which APPID has not yet made either a final or best-effort determination of the associated application.

Upon Best-Effort Application Identification

When a best-effort application determination is made, AACL does not apply any AACL term actions configured for that flow. There are a number of reasons for this, one being that the action itself (such as "discard") could make a final application determination impossible. Instead, AACL or L-PDF tracks the flow and accepts all packets for that flow until a final determination is made, at which time the normal AACL or L-PDFL actions are fully applied to the flow.

While Application Identification Is on a Best-Effort Basis

During the time that APPID identification of the application associated with a given flow is on a best-effort basis, the flow does not contribute to any per-subscriber or per-application statistics collection.

The output of the following operational mode commands includes flows for which APPID has only made a best-effort determination of the associated application:

- **show services local-policy-decision-function flows** (interface *interface-name* | subscriber *subscriber-name*)
- **show services application-aware-access-list flows** (interface *interface-name* | subscriber *subscriber-name*)

In the command output, the **Action** field displays "accept" and the **Application** or **Application group** field displays "unknown" for a flow for which APPID has only made a best-effort determination of the associated application.

If a Flow Ends Before an Application Identification Is Made

If a flow ends before APPID has made either a final or a best-effort application identification, AACL or L-PDF uses the "unknown" application ID as a final determination and performs any necessary collection, aggregation, and reporting of statistics based on that Layer 7 application. In particular, if the **count** AACL term action is configured for the "application-group-any" application, then the statistics for that flow will be collected and aggregated against the count bucket type, and reported as such.

If a Flow Ends While Application Identification on a Best-Effort Basis

If a flow ends while the application identification is on a best-effort basis, AACL or L-PDF uses that best-effort determination as a final determination. AACL or L-PDF performs

any necessary collection, aggregation, and reporting of statistics based on that Layer 7 application. In particular, if the **count** AACL term action is configured for that Layer 7 application, then the statistics for the flow will be collected and aggregated against the AACL or L-PDF statistics. However, in the case of nested applications, AACL and L-PDF will not consider the best-effort determination as final and the nested application will be reported as an unknown application.

Related Documentation

- [Configuring AACL Rules on page 661](#)
- [Configuring Statistics Profiles on page 702](#)
- [aACL-fields on page 1539](#)
- [aACL-statistics-profile on page 1540](#)
- [rule on page 1575](#)
- [services on page 1582](#)
- [term on page 1588](#)
- [then on page 1589](#)

Configuring AACL Rules

To configure an AACL rule, include the **rule** *rule-name* statement at the **[edit services aACL]** hierarchy level:

```
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application-group-any;
      application-groups [ application-group-names ];
      applications [ application-names ];
      destination-address address <any-unicast>;
      destination-address-range low minimum-value high maximum-value;
      destination-prefix-list list-name;
      nested-applications [ nested-application-names ];
      nested-application-unknown;
      source-address address <any-unicast>;
      source-address-range low minimum-value high maximum-value;
      source-prefix-list list-name;
    }
    then {
      (accept | discard);
      count (application | application-group | application-group-any | nested-application
        | none);
      forwarding-class class-name;
      policer policer-name;
    }
  }
}
```

Each AACL rule consists of a set of terms, similar to a filter configured at the **[edit firewall]** hierarchy level. A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded.
- **then** statement—Specifies the actions and action modifiers to be performed by the router software.

The following sections explain how to configure the components of ACL rules:

- [Configuring Match Direction for ACL Rules on page 662](#)
- [Configuring Match Conditions in ACL Rules on page 662](#)
- [Configuring Actions in ACL Rules on page 664](#)
- [Logging ACL Flows Based on Application on page 665](#)

Configuring Match Direction for ACL Rules

Each rule must include a **match-direction** statement that specifies the direction in which the rule match is applied. To configure where the match is applied, include the **match-direction** statement at the **[edit services acl rule *rule-name*]** hierarchy level:

```
match-direction (input | output | input-output);
```

If you configure **match-direction input-output**, bidirectional rule creation is allowed.

The match direction is used with respect to the traffic flow through the services PIC or DPC. When a packet is sent to the PIC or DPC, direction information is carried along with it.

With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.

With a next-hop service set, packet direction is determined by the interface used to route the packet to the services PIC or DPC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the PIC or DPC, the packet direction is output. For more information on inside and outside interfaces, see [“Configuring Service Sets to be Applied to Services Interfaces” on page 31](#).

On the PIC or DPC, a flow lookup is performed. If no flow is found, rule processing is performed. All rules in the service set are considered. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that matches the packet direction are considered.

Configuring Match Conditions in ACL Rules

To configure ACL match conditions, include the **from** statement at the **[edit services acl rule *rule-name* term *term-name*]** hierarchy level:

```
from {  
  application-group-any;  
  application-groups [ application-group-names ];  
  applications [ application-names ];  
  destination-address address <any-unicast>;  
  destination-address-range low minimum-value high maximum-value;  
  destination-prefix-list list-name;
```



```

nested-applications [ nested-application-names ];
nested-application-unknown
source-address address <any-unicast>;
source-address-range low minimum-value high maximum-value;
source-prefix-list list-name;
}

```

IPv4 and IPv6 source and destination addresses are supported. You can use either the source address or the destination address as a match condition, in the same way that you configure a firewall filter; for more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Alternatively, you can specify a list of source or destination prefixes by configuring the **prefix-list** statement at the **[edit policy-options]** hierarchy level and then including either the **destination-prefix-list** or the **source-prefix-list** statement in the ACL rule. For an example, see “[Example: Configuring ACL Rules](#)” on page 666.

If you omit the **from** term, the ACL rule accepts all traffic and the default protocol handlers take effect:

- User Datagram Protocol (UDP), Transmission Control Protocol (TCP), and Internet Control Message Protocol (ICMP) create a bidirectional flow with a predicted reverse flow.
- IP creates a unidirectional flow.

You can also include application and application group definitions you have configured at the **[edit services application-identification]** hierarchy level; for more information, see the topics in *Application Identification*.

- To apply one or more specific application protocol definitions, include the **applications** statement at the **[edit services aacl rule *rule-name* term *term-name* from]** hierarchy level.
- To apply one or more sets of application group definitions you have defined, include the **application-groups** statement at the **[edit services aacl rule *rule-name* term *term-name* from]** hierarchy level.



NOTE: If you include one of the statements that specifies application protocols, the router derives port and protocol information from the corresponding configuration at the **[edit services application-identification]** hierarchy level; you cannot specify these properties as match conditions.

- To consider any application group defined in the database as a match, include the **application-group-any** statement at the **[edit services aacl rule *rule-name* term *term-name* from]** hierarchy level.
- To consider any nested application defined in the database a match, include the **nested-applications** statement at the **[edit services aacl rule *rule-name* term *term-name* from]** hierarchy level. Nested applications are protocols that run on a parent application. For example, if the Facebook application runs on the parent application `junos:http`, the nested application will be `junos:http:facebook`.

Configuring Actions in ACL Rules

To configure ACL actions, include the **then** statement at the **[edit services aac rule rule-name term term-name]** hierarchy level:

```
then {  
  (accept | discard);  
  (count (application | application-group | application-group-any | nested-application |  
    none) | forwarding-class class-name);  
}
```

You must include one of the following actions:

- **accept**—The packet is accepted and sent on to its destination.
- **discard**—The packet is not accepted and is not processed further.

When you select **accept** as the action, you can optionally configure one or both of the following action modifiers. No action modifiers are allowed with the **discard** action.

- **count (application | application-group | application-group-any | nested-application | none)**—For all accepted packets that match the rules, record a packet count using ACL statistics practices. You can specify one of the following options; there is no default setting:
 - **application**—Count the application that matched in the **from** clause.
 - **application-group**—Count the application group that matched in the **from** clause.
 - **application-group-any**—Count all application groups that match **from application-group-any** under the **any** group name.
 - **nested-application**—Count all nested applications that matched in the **from** clause.
 - **none**—Same as not specifying **count** as an action.



NOTE:

- When a session closes before APPID has identified nested applications, the session is treated as a best-effort session and ACL does not get the nested application information. In such cases, nested applications will be reported as unknown applications.
- During the time that the application identification (APPID) feature has not yet made a final determination of the application associated with a given flow, the flow does not contribute to any per-subscriber or per-application statistics collection. For more information, see [“Best-Effort Application Identification of DPI-Serviced Flows” on page 658](#).

-
- **forwarding-class class-name**—Specify the packets’ forwarding-class name.

You can optionally include a **policer** that has been specified at the **[edit firewall]** hierarchy level. Only the bit-rate and burst-size properties specified for the policer are applied in the ACL rule set. The only action application when a policer is configured is **discard**. For

more information on policer definitions, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Logging AACL Flows Based on Application

You can now log AACL flows based on application. You can select a specific application or request information on unknown applications.

You can now configure AACL rules to match unknown applications. All existing actions that can apply to recognized applications can also apply to unknown applications. You can use the following statements at the **[edit services aacl rule rule-name term term-name from]** hierarchy level:

- application-group-any
- application-groups
- application-unknown
- applications
- nested-application-unknown
- nested-applications

The addition of matching “application unknown” enables the specific logging of the input flows associated with applications that cannot be identified. Because logging is triggered by an input event, you must specify **match-direction** as **input-output** or **input**.

To configure logging of flows for AACL, include the **match-direction input** or **match-direction input-output** statement at the **[edit services aacl rule rule-name]** hierarchy level, include an **applications** or **application-unknown** statement at the **[edit services aacl rule rule-name term term-name from]** hierarchy level, and include only one **log** statement at the **[edit services aacl rule rule-name term term-name then]** hierarchy level. The log statements can include any of the following options:

- session-start
- session-end
- session-start-end-no-stats
- session-start-interim-end
- session-interim-end
- session-end

Related Documentation

- [AACL Overview on page 657](#)
- [Configuring AACL Rule Sets on page 666](#)
- [Configuring Logging of AACL Flows on page 667](#)
- [Example: Configuring AACL Rules on page 666](#)

Example: Configuring AACL Rules

The following example shows an AACL configuration containing a rule with three terms using a variety of match conditions and actions:

```
[edit services aacl]
rule aacl-test {
  match-direction input;
  term term1 {
    from {
      source-address 10.0.1.1
      application test1;
    }
    then {
      accept;
    }
  }
  term term2 {
    from {
      source-address {
        any-unicast;
      }
      application test1;
    }
    then {
      discard;
    }
  }
  term term3 {
    from {
      source-address {
        any-unicast;
      }
      application test1 test2;
    }
    then {
      accept;
      count application;
    }
  }
}
```

- Related Documentation**
- [AACL Overview on page 657](#)
 - [Configuring AACL Rules on page 661](#)

Configuring AACL Rule Sets

The **rule-set** statement defines a collection of AACL rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by

including the **rule-set** statement at the **[edit services aacl]** hierarchy level with a **rule** statement for each rule:

```
rule-set rule-set-name {
  rule rule-name;
}
```

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

Related Documentation

- [AACL Overview on page 657](#)
- [Configuring AACL Rules on page 661](#)
- [Configuring Logging of AACL Flows on page 667](#)
- [Example: Configuring AACL Rules on page 666](#)

Configuring Logging of AACL Flows

You can configure logging of AACL flows for a given application or for all unknown applications using AACL rules. You must set **match-direction** to **input** or **input-output** for logging to occur.

1. Create a rule and term.

```
user@host# edit services aacl rule rule-name term term-name
```

2. Specify selection of an application.

```
[edit services aacl rule rule-name term term-name]
user@host# set from applications application-name]
```

OR

Specify selection of all unknown applications.

```
[edit services aacl rule <variable>rule-name</variable> term
<variable>term-name</variable>]
set from application-unknown
```

3. In the **then** statement, specify logging of input flow.

```
[edit services aacl rule rule-name term term-name]
user@host# set then log input-flows]
```

Example—Configuration of Logging of Input Flows for Unknown Applications

```
[edit services aacl rule aacL_rule5]
match-direction input-output;
term t0 {
  from {
    application-unknown;
  }
  then {
    count application;
```

```
        log input-flow;
        accept;
    }
}
```

Example—Setup of a Specific Log File

The following example shows how to direct the aacL flow log to a file other than the default syslog file on the Routing Engine file system.

```
[edit system syslog]
file aacL_log {
    external any;
    match aacL-flow-log;
}
```

Related Documentation

- [AACL Overview on page 657](#)
- [Configuring AACL Rules on page 661](#)
- [Configuring AACL Rule Sets on page 666](#)
- [Example: Configuring AACL Rules on page 666](#)

Grouping Applications Together Using APPID

- [APPID Overview on page 669](#)
- [Best-Effort Application Identification of DPI-Serviced Flows on page 671](#)
- [Defining an Application Identification on page 674](#)
- [Configuring APPID Rules on page 676](#)
- [Using Stateful Firewall Rules to Identify Data Sessions on page 677](#)
- [Configuring Application Profiles on page 679](#)
- [Configuring Application Groups on page 680](#)
- [Application Identification for Nested Applications on page 681](#)
- [Disabling Application Identification for Nested Applications on page 682](#)
- [Configuring Global APPID Properties on page 683](#)
- [Configuring APPID Support for Heuristics on page 684](#)
- [Configuring APPID Support for Unidirectional Traffic on page 685](#)
- [Configuring Automatic Download of Application Package Updates on page 686](#)
- [Tracing APPID Operations on page 686](#)
- [Examples: Configuring Application Identification Properties on page 688](#)

APPID Overview



NOTE: Starting with Junos OS Release 12.1, all interface-style services are supported for dynamic Point-to-Point Protocol over Ethernet (PPPoE) subscribers on all MX Series routers with modular Modular Port Concentrators (MPCs).

The APPID feature identifies applications as constituents of application groups in TCP/UDP/ICMP traffic. It is supported on MX Series routers equipped with Multiservices DPCs and on M120 or M320 routers equipped with Multiservices 400 PICs and Aggregated Multiservices (AMS) PICs. Aggregated Multiservices PICs (ams- interfaces) enable multiple ms- interfaces to be grouped together in a single bundle and cause the traffic destined for this AMS group to be distributed over the member services PICs of the group.

Junos OS Trio chipsets enable the calculation of a symmetric hash for the forward and reverse flows, and support a microcode map in the forwarding plane. This capability enables load-balancing of traffic across various services PICs in an AMS group. Starting with Junos OS Release 12.1, ams- interfaces enable an N:1 redundancy mechanism to cluster together N number of ms- interfaces in an AMS group that supports load sharing.



NOTE: For ams- interfaces and rms- interfaces, the statistics data in the bulk statistics file is collected using the reports received from the MS PICs. For the ams- interfaces, the retrieval and storage of statistics is not possible because of multiple PICs containing statistics data for the same subscriber. For interfaces in an AMS group, statistics data from different MS PICs in the AMS group are collected and aggregated on the Routing Engine where a timer control is activated and the data is saved in the bulkstats file based on this timer. This method of collection causes the statistics data in the bulkstats file to be displayed with a small delay period.

To configure APPID, include statements at the **[edit services application-identification]** hierarchy level to specify parameter values for defining applications, enable or disable application rules, and gather the applications and rules into groups.

The following are related operational commands:

- **show/clear application-identification application-system-cache**
- **show/clear application-identification counters**

For more information on the CLI configuration, see the *Application Identification*. For more information on the operational commands, see the [CLI Explorer](#).



NOTE: Because the extension-provider package framework lacks aggressive constraint checks, you should not set the `policy-db-size` statement at the `[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]` hierarchy level to a high value. For Junos Application Aware (previously known as Dynamic Application Awareness) configurations, the recommended values for the extension-provider package options at this hierarchy level are as follows:

- `control-cores = 1`
- `data-cores = 7`
- `object-cache-size = 1280` (for Multiservices 400 PIC and Multiservices DPC)
- `policy-db-size = 200`
- Include these package values: `jservices-idp`, `jservices-appid`, `jservices-llpdf`, `jservices-aac1`

For more information about this configuration, see the following topics in the *SDK Applications Configuration Guide and Command Reference*:

- *Configuring Control and Data Cores*
- *Configuring Memory Settings*
- *Configuring Packages on the PIC*



NOTE: In the export version of JUNOS, signature download is not expected to work for AppID and IDP features in Junos Application Aware. In order to make it work, you must additionally install the Crypto Software Suite.

Related Documentation

- [Defining an Application Identification on page 674](#)
- [Configuring APPID Rules on page 676](#)
- [Application Identification for Nested Applications on page 681](#)
- [Configuring Global APPID Properties on page 683](#)
- [Examples: Configuring Application Identification Properties on page 688](#)
- [\[edit services application-identification\] Hierarchy Level on page 1295](#)

Best-Effort Application Identification of DPI-Serviced Flows

This topic describes the following information:

- [Features that Support Application-Level Filtering on page 672](#)
- [Best-Effort Application Determination on page 672](#)
- [APPID, AACL, and L-PDF Processing in Preconvergence Scenarios on page 672](#)

Features that Support Application-Level Filtering

On MX Series routers equipped with Multiservices DPCs and M120 or M320 routers equipped with Multiservices 400 PICs, Intrusion Detection and Prevention (IDP) is accomplished by Deep Packet Inspection (DPI) of TCP, UDP, and ICMP flows. The application identification (APPID) feature defines applications as members of application groups in TCP/UDP/ICMP traffic. IDP depends on APPID for identification and detection of some Layer 7 applications.

The application-aware access list (AACL) service uses application names and groups as matching criteria for filtering traffic. The service defines the applications and application groups for which statistics are collected for a specific user or interface.

The local policy decision function (L-PDF) enables you to configure properties for statistics output. L-PDF supports a process that regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

Best-Effort Application Determination

Typically, APPID conclusively determines the Layer 7 application associated with a given DPI-serviced flow. In these cases, the application identification is final. Occasionally, APPID is only able to make an initial, inconclusive determination of the Layer 7 application associated with a given flow. This is referred to as a "best-effort" application identification. In such cases, the APPID process continues processing packets on that flow and might subsequently make a conclusive determination of the application associated with that flow. In some cases of best-effort application identification, the flow ends before a final application determination can be made.

APPID, AACL, and L-PDF Processing in Preconvergence Scenarios

The following sections describe APPID, AACL, and L-PDF processing in various stages of application identification for a DPI-serviced flow of TCP/UDP/ICMP traffic.

- [Prior to a Final or Best-Effort Application Identification on page 672](#)
- [Upon Best-Effort Application Identification on page 673](#)
- [While Application Identification Is on a Best-Effort Basis on page 673](#)
- [If a Flow Ends Before an Application Identification Is Made on page 673](#)
- [If a Flow Ends While Application Identification on a Best-Effort Basis on page 673](#)

Prior to a Final or Best-Effort Application Identification

During the time that APPID has not yet made either a final or best-effort determination of the application associated with a given flow, the flow does not contribute to any per-subscriber or per-application statistics collection.

The output of the following operational mode commands includes flows for which APPID has not yet made either a final or best-effort determination of the associated application:

- **show services local-policy-decision-function flows** (interface *interface-name* | subscriber *subscriber-name*)
- **show services application-aware-access-list flows** (interface *interface-name* | subscriber *subscriber-name*)

In the command output, the **Action** field displays "accept" and the **Application** or **Application group** field displays "unknown" for a flow for which APPID has not yet made either a final or best-effort determination of the associated application.

Upon Best-Effort Application Identification

When a best-effort application determination is made, AACL does not apply any AACL term actions configured for that flow. There are a number of reasons for this, one being that the action itself (such as "discard") could make a final application determination impossible. Instead, AACL or L-PDF tracks the flow and accepts all packets for that flow until a final determination is made, at which time the normal AACL or L-PDFL actions are fully applied to the flow.

While Application Identification Is on a Best-Effort Basis

During the time that APPID identification of the application associated with a given flow is on a best-effort basis, the flow does not contribute to any per-subscriber or per-application statistics collection.

The output of the following operational mode commands includes flows for which APPID has only made a best-effort determination of the associated application:

- **show services local-policy-decision-function flows** (interface *interface-name* | subscriber *subscriber-name*)
- **show services application-aware-access-list flows** (interface *interface-name* | subscriber *subscriber-name*)

In the command output, the **Action** field displays "accept" and the **Application** or **Application group** field displays "unknown" for a flow for which APPID has only made a best-effort determination of the associated application.

If a Flow Ends Before an Application Identification Is Made

If a flow ends before APPID has made either a final or a best-effort application identification, AACL or L-PDF uses the "unknown" application ID as a final determination and performs any necessary collection, aggregation, and reporting of statistics based on that Layer 7 application. In particular, if the **count** AACL term action is configured for the "application-group-any" application, then the statistics for that flow will be collected and aggregated against the count bucket type, and reported as such.

If a Flow Ends While Application Identification on a Best-Effort Basis

If a flow ends while the application identification is on a best-effort basis, AACL or L-PDF uses that best-effort determination as a final determination. AACL or L-PDF performs

any necessary collection, aggregation, and reporting of statistics based on that Layer 7 application. In particular, if the **count** AACL term action is configured for that Layer 7 application, then the statistics for the flow will be collected and aggregated against the AACL or L-PDF statistics. However, in the case of nested applications, AACL and L-PDF will not consider the best-effort determination as final and the nested application will be reported as an unknown application.

**Related
Documentation**

- [Configuring AACL Rules on page 661](#)
- [Configuring Statistics Profiles on page 702](#)
- [aACL-fields on page 1539](#)
- [aACL-statistics-profile on page 1540](#)
- [rule on page 1575](#)
- [services on page 1582](#)
- [term on page 1588](#)
- [then on page 1589](#)

Defining an Application Identification

To configure a specific IP address or port-based application identification, include the **application** *application-name* statement at the **[edit services application-identification]** hierarchy level:

```
application application-name {  
  disable;  
  idle-timeout seconds;  
  index number;  
  session-timeout seconds;  
  type type;  
  type-of-service service-type;  
  port-mapping {  
    port-range {  
      tcp [ ports-and-port-ranges ];  
      udp [ ports-and-port-ranges ];  
    }  
    disable;  
  }  
}
```

You can include the following general properties in the configuration:

- **application**—Application name, a required statement; maximum 31 characters. Predefined applications have the prefix **junos-** to avoid conflict with user-defined ones.
- **idle-timeout**—Amount of time that a session remains idle before it is deleted.
- **index**—Application index number in the range from 1 through 65,534, with integers 1 through 1024 reserved for predefined applications.
- **session-timeout**—Lifetime of a session.

- **type**—Well known applications, such as HTTP or FTP.
- **type-of-service**—Type of service, defined by service objective. There is no default value; options are **maximize-reliability**, **maximize-throughput**, **minimize-delay**, and **minimize-monetary-cost**.
- **disable**—Disable this application definition in the APPID service.



NOTE: You can also specify session and idle timeout values globally for a Multiservices interface by including the following statements at the [edit interfaces *interface-name* services-options] hierarchy level:

- **inactivity-non-tcp-timeout**—Inactivity timeout period for non-TCP established sessions.
- **inactivity-tcp-timeout**—Inactivity timeout period for TCP established sessions.
- **session-timeout**—Lifetime of a session.
- **disable-global-timeout-override**—Disallow overriding a global inactivity or session timeout.

You can include the following port-mapping properties at the [edit services application-identification port-mapping] hierarchy level:

- **port-range**—TCP or UDP port number or numeric range, entered as [*minimum-value* – *maximum-value*]. For port-mapping configurations, this entry is required if the parent node exists.
- **disable**—Disable port-mapping properties for this application.



NOTE: For applications with signatures for both client-to-server and server-to-client directions, the APPID for Junos Application Aware (previously known as Dynamic Application Awareness) must accept the data packets in both directions on the same session to complete the identification process.

For a configuration example, see “[Examples: Configuring Application Identification Properties](#)” on page 688.

Related Documentation

- [APPID Overview on page 669](#)
- [Configuring APPID Rules on page 676](#)
- [Using Stateful Firewall Rules to Identify Data Sessions on page 677](#)
- [Configuring Application Profiles on page 679](#)
- [Configuring Application Groups on page 680](#)
- [Tracing APPID Operations on page 686](#)

- [\[edit services application-identification\] Hierarchy Level on page 1295](#)

Configuring APPID Rules

This configuration specifies the properties for identifying an application for which a source or destination IP address and port is used for a known application, without the requirement of an application signature. For example, the Session Initiation Protocol (SIP) server initiates a session from its identified port, 5060. You can therefore specify the SIP server IP address and port 5060 in the port mapping configuration for the SIP application. The advantage of using this method is to provide efficiency and accuracy of application identification for your network.

To configure application rule properties, include the **rule** statement at the [\[edit services application-identification\]](#) hierarchy level:

```
rule rule-name {  
  address address-name {  
    destination {  
      ip address</prefix-length>;  
      port-range {  
        tcp [ ports-and-port-ranges ];  
        udp [ ports-and-port-ranges ];  
      }  
    }  
    source {  
      ip address</prefix-length>;  
      port-range {  
        tcp [ ports-and-port-ranges ];  
        udp [ ports-and-port-ranges ];  
      }  
    }  
    order number;  
  }  
  application application-name;  
  disable;  
}
```

You can include the following application rule properties:

- **address**—Address properties for APPID rule processing. This statement is mandatory; you must specify either destination or source properties.
- **destination**—Destination address and port information. The **ip** statement defines the IP address and netmask (IPv4 only), and the **port-range** statement defines the TCP or UDP port number or numeric range, entered as **[*minimum-value* – *maximum-value*]**.
- **source**—Source address and port information. The **ip** statement defines the IP address and netmask (IPv4 only), and the **port-range** statement defines the TCP or UDP port number or numeric range, entered as **[*minimum-value* – *maximum-value*]**.
- **order**—Application matching priority. For address configurations, the order number resolves the conflict when multiple address entries are matched for a specific session;

the lower the number, the higher the priority. This statement is mandatory and must contain a unique value.

- **application**—Name of the application to be included in the rule.
- **disable**—Disable processing for this application rule.

The **rule-set** statement defines a collection of APPID rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services application-identification]** hierarchy level with a **rule** statement for each rule:

```
rule-set rule-set-name {
  rule application-rule-name;
}
```

Related Documentation

- [APPID Overview on page 669](#)
- [Defining an Application Identification on page 674](#)
- [Using Stateful Firewall Rules to Identify Data Sessions on page 677](#)
- [Configuring Application Profiles on page 679](#)
- [Configuring Application Groups on page 680](#)
- [Examples: Configuring Application Identification Properties on page 688](#)

Using Stateful Firewall Rules to Identify Data Sessions

The APPID configuration properties enable the Junos OS to detect applications based on signatures, ports, and addresses. For signature-based detection, most of the protocol control sessions are identified, but data sessions are not identified. For example, APPID identifies FTP connections to port 21 (FTP control sessions); however, FTP can open child/data sessions to transfer files and data. These sessions are not identified by signature-based APPID because they do not have well-defined signatures.

Application-level gateways (ALGs) configured using stateful firewall rules can assist APPID in identifying these data sessions. These sessions include file and video transfers that are heavy consumers of bandwidth, so a mechanism for policing and classifying this traffic effectively is a useful tool. In addition to FTP, this mechanism applies to TFTP and RTSP traffic.

To incorporate the stateful firewall rules into Junos Application Aware (previously known as Dynamic Application Awareness for Junos OS) sessions, include the following configurations:

1. Include the stateful firewall package at the **[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]** hierarchy level:


```
package jservices-sfw;
```
2. Define two stateful firewall rules as shown in the following example, one to identify the appropriate ALGs for FTP, TFTP, or RTSP traffic and the other to allow all traffic:



NOTE: Session Initiation Protocol (SIP) is already covered by APPID and the SIP ALG is not supported by stateful firewall, hence a SIP configuration is not needed.

```
[edit services]
stateful-firewall {
  rule rule1 {
    match-direction input-output;
    term term1 {
      from {
        applications [ junos-ftp junos-tftp junos-rtsp ];
      }
      then {
        accept;
      }
    }
  }
  rule rule2 {
    match-direction input-output;
    term term1 {
      then {
        accept;
      }
    }
  }
  rule-set rs1 {
    rule rule1;
    rule rule2;
  }
}
```



NOTE: The existing AACL and L-PDF operational mode commands should report the new applications when they are identified.

3. Attach the stateful firewall rule set to a service set, as shown in the following example:

```
service-set test-chaining {
  application-identification-profile add-based;
  stateful-firewall-rule-sets rs1;
  idp-profile idp1;
  aacl-rules rule1;
  interface-service {
    service-interface ms-2/0/0.0;
  }
}
```

4. Include *no-drop* settings for stateful firewall and TCP, as needed.

Stateful firewall processing drops packets in a number of scenarios:

- TCP sessions do not start with a SYN flag. (This prevents sessions from resuming; otherwise, when the PIC starts for the first time, all existing TCP sessions in flight will be dropped).
- If the TCP tracker detects SYN but no SYN/ACK or only an ACK, then the ACK is dropped. There are a number of similar checks to verify the TCP connection, window checks, and so forth.
- TCP checks for stateful firewall are aggressive when ALGs are run. It is not possible to ignore TCP errors when an ALG is run on a session.
- If an ALG detects malformed packets (for example, if the FTP PORT command is not RFC-compliant), it drops packets. If an ALG is not able to allocate resources, it drops packets.

You can include the settings shown in the following example to assist in controlling these packet drops:

```
[edit interfaces]
ms-1/2/0 {
  services-options {
    ignore-errors {
      tcp;
      alg;
    }
  }
}
```

The **tcp** statement mediates the first two issues listed, with reference to TCP SYN detection. The **alg** statement handles the fourth issue. ALGs require strict TCP processing, which cannot be relaxed.

**Related
Documentation**

- [APPID Overview on page 669](#)
- [Defining an Application Identification on page 674](#)
- [Application Identification for Nested Applications on page 681](#)
- [Configuring Global APPID Properties on page 683](#)
- [Tracing APPID Operations on page 686](#)

Configuring Application Profiles

You can define an application profile for use in a service set. The profile consists of one or more rule sets, but only one profile can be included per service set.

To specify the application profile constituents, include the **profile** statement at the **[edit services application-identification]** hierarchy level:

```
profile profile-name {
  [ rule-set rule-set-name ];
}
```

You assign a profile name and include one or more predefined rule sets. For more information on rule sets, see [“Configuring APPID Rules” on page 676](#). You can then include the profile in a service-set definition:

```
[edit services]
service-set service-set-name {
  profile profile-name;
}
```

The definitions specific to Junos Application Aware (previously known as Dynamic Application Awareness) include the APPID and IDP profiles and the ACL rule set. For more information on service sets, see *Service Set Properties*.

- Related Documentation**
- [APPID Overview on page 669](#)
 - [Defining an Application Identification on page 674](#)
 - [Configuring Application Groups on page 680](#)
 - [Configuring Global APPID Properties on page 683](#)

Configuring Application Groups

You can define an application group to process a number of applications or subgroups at the same time. To configure application group properties, include the **application-group** statement at the **[edit services application-identification]** hierarchy level:

```
application-group group-name {
  application-groups {
    application-group-name;
  }
  applications {
    application-name;
  }
  index number;
  disable;
}
```

You can include the following application group properties:

- **applications**—List of applications to include in this application group. The **name** statement is mandatory and must include at least one entry.
- **application-groups**—List of application groups to include in a larger application group. The **name** statement is mandatory and must include at least one entry.
- **index**—Application group index number in the range from 1 through 65,534. This mandatory value must be unique.
- **disable**—Disable processing for this application group.

- Related Documentation**
- [Defining an Application Identification on page 674](#)
 - [Configuring APPID Rules on page 676](#)
 - [Configuring Application Profiles on page 679](#)

- [Configuring Global APPID Properties on page 683](#)
- [Examples: Configuring Application Identification Properties on page 688](#)

Application Identification for Nested Applications

The application identification feature is used by intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports. *Nested applications* are protocols running over the parent application. For example, both Facebook and Yahoo Messenger can run over HTTP, but there is a need to identify them as two different applications. To do this, the application layer is split into two layers: Layer 7 applications and Layer 7 protocols.

The predefined application signatures included with Junos OS have been created to detect the Layer 7 nested applications. Predefined application signatures can be used in attack objects.

To configure nested application properties, include the **nested-application** statement at the **[edit services application-identification]** hierarchy level:

```
nested-application name {
  index number;
  protocol protocol;
  signature name {
    chain-order;
    maximum-transactions number;
    member name {
      context (http-header-content-type | http-header-host | http-url-parsed |
        http-url-parsed-param-parsed);
      direction (any | client-to-server | server-to-client);
      pattern dfa-pattern;
    }
    order number;
  }
  type type;
}
```

You can include the following application rule properties:

- **chain-order**—Signatures can contain multiple members. If the chain order feature is on, those members are read in order. The default for this option is no chain order. If a signature contains only one member, this option is ignored.
- **context**—Define a service specific context. The options are **http-header-content-type**, **http-header-host**, **http-url-parsed**, **http-url-parsed-param-parsed**. This statement is mandatory.
- **direction**—The connection direction of the packets to apply pattern matching. The options are **client-to-server**, **server-to-client**, or **any**. This statement is mandatory.
- **index**—A number that is a one-to-one mapping to the application name that is used to ensure that each signature definition is unique. The index range for predefined

applications is 1 through 32767. The index range for custom applications and custom nested applications is 32768 through 65534.

- **maximum transactions**—The maximum number of transactions that should occur before a match is made. This statement is mandatory.
- **member**—Define a member name for a custom nested application signature definition. Custom definitions can contain multiple members that define attributes for an application.
- **order**—Define application matching priority. For address configurations, the order number resolves the conflict when multiple address entries are matched for a specific session. The lower number has higher priority. This statement is mandatory.
- **pattern**—Define an attack pattern to be detected. This statement is mandatory.
- **protocol**—The protocol that will be monitored to identify nested applications. The value **http** is supported. This statement is mandatory.
- **signature**—Name of the custom nested application signature definition. Must be a unique name with a maximum length of 32 characters. This statement is mandatory.
- **type**—Well-known application name for this application definition, such as Facebook or Kazza. This application name must be unique with a maximum length of 32 characters. This statement is mandatory.

Related Documentation

- [APPID Overview on page 669](#)
- [Defining an Application Identification on page 674](#)
- [Disabling Application Identification for Nested Applications on page 682](#)
- [Configuring Global APPID Properties on page 683](#)
- [Tracing APPID Operations on page 686](#)

Disabling Application Identification for Nested Applications

Sometimes there is a need to identify multiple different applications running on the same Layer 7 protocols. For example, both Facebook and Yahoo Messenger can run over HTTP, but there is a need to identify them as two different applications. To do this, the application layer is split into two layers: Layer 7 applications and Layer 7 protocols. Application identification for nested applications is turned on by default. You can manually turn it off by using the CLI.

To disable nested application identification:

- Set the **no-nested-application** statement.

```
[edit services application-identification nested-application-settings]  
user@host# no-nested-application
```

To verify the configuration, issue the **show services application-identification nested-application-settings** command.

To reenable nested application identification:

- Delete the **no-nested-application** statement.

```
[edit services application-identification nested-application-settings]
user@host# delete services application-identification nested-application-settings
no-nested-application
```

If you are finished configuring the device, commit the configuration.

**Related
Documentation**

- [APPID Overview on page 669](#)
- [Application Identification for Nested Applications on page 681](#)

Configuring Global APPID Properties

You can define additional properties that apply on a global basis to APPID processing and are not part of a specific application, group, rule, or profile definition. To configure these global APPID properties, include the following statements at the **[edit services application-identification]** hierarchy level:

```
application-identification {
  application-system-cache-timeout seconds;
  max-checked-bytes bytes;
  min-checked-bytes bytes;
  nested-application name
  nested-application-settings
  no-application-identification
  no-application-system-cache;
  no-clear-application-system-cache;
  no-protocol-method;
  no-signature-based;
  signature-method-all-ports;
}
```

The global application properties have the following effect:

- **application-system-cache-timeout**—Lifetime for system cache entries, in seconds.
- **max-checked-bytes**—The maximum number of bytes to be inspected in APPID processing, in the range from 0 through 100,000 bytes.
- **min-checked-bytes**—The minimum number of bytes to be inspected in APPID processing, in the range from 0 through 2000 bytes.
- **nested-application**—Configure a custom nested application definition for the desired application name that will be used by the system to identify the nested application as it passes through the device. For more information see [nested-application](#).
- **nested-application-settings**—Configure nested application options for application identification services. For more information see [nested-application-settings](#).
- **no-application-identification**—Disable all application identification methods.
- **no-application-system-cache**—Disable storing application identification results in the application system cache.

- **no-clear-application-system-cache**—Disable clearing the application system cache.
- **no-protocol-method**—Disable the protocol-based application identification method, which is enabled by default.
- **no-signature-based**—Disable the signature-based application identification method.
- **signature-method-all-ports**—Run signature matching on all traffic.

Related Documentation

- [APPID Overview on page 669](#)
- [Defining an Application Identification on page 674](#)
- [Application Identification for Nested Applications on page 681](#)
- [Disabling Application Identification for Nested Applications on page 682](#)
- [Tracing APPID Operations on page 686](#)
- [Examples: Configuring Application Identification Properties on page 688](#)

Configuring APPID Support for Heuristics

Heuristics methodology provides a mechanism for identifying encrypted data packets in point-to-point applications. These packets are not normally detected by the existing application signatures.

To enable APPID to employ heuristics in traffic identification:

1. Include the **enable-heuristics** statement:

```
[edit services application-identification]
user@host# enable-heuristics
```

The **show services application-identification counter** operational command includes additional output fields that report the number of encrypted sessions.



NOTE: When you enable heuristics, performance and scaling values might be negatively affected. This mechanism assists the APPID module in identifying encrypted traffic, but only if the identifications are supported by the current signature package.

Related Documentation

- [APPID Overview on page 669](#)
- [Defining an Application Identification on page 674](#)
- [Application Identification for Nested Applications on page 681](#)
- [Configuring Global APPID Properties on page 683](#)
- [Configuring APPID Support for Unidirectional Traffic on page 685](#)
- [Examples: Configuring Application Identification Properties on page 688](#)

Configuring APPID Support for Unidirectional Traffic

With asymmetrical routing, a networking device sees only one side of the network sessions, either from client to server or from server to client. Additional functionality is required to support application identification with unidirectional traffic. This addition enables a session for a specified service set to support an asymmetrical routing environment, and allows complete application matches using existing application signatures for traffic in the client-to-server direction only.

To enable APPID to support application matching on unidirectional traffic:

1. Include the **support-uni-directional-traffic** statement:

```
[edit services service-set service-set-name service-set-options]
user@host# support-uni-directional-traffic
```

This enables the session belonging to the specified service set to support the asymmetrical routing environment. The APPID module then reports complete matches for the unidirectional traffic.

2. Include the **enable-asymmetric-traffic-processing** statement:

```
[edit services service-set service-set-name service-set-options]
user@host# enable-asymmetric-traffic-processing
```

This enables the framework and plug-in to handle unidirectional traffic at a service-set level.

When you enable these settings, APPID treats unidirectional TCP traffic like a UDP connection. UDP traffic itself does not receive any special treatment because the service PIC cannot determine whether UDP traffic is unidirectional or bidirectional. The settings do not affect processing of sessions created with bidirectional traffic.

If the traffic includes both unidirectional and bidirectional sessions, the APPID module uses heuristics to decide whether to change the reporting logic.



NOTE: This feature does not change the processing for any services except APPID. However, other services, including stateful firewall, AACL, and IDP, can process unidirectional traffic in a limited manner.

Related Documentation

- [APPID Overview on page 669](#)
- [Defining an Application Identification on page 674](#)
- [Application Identification for Nested Applications on page 681](#)
- [Configuring Global APPID Properties on page 683](#)
- [Configuring APPID Support for Heuristics on page 684](#)
- [Examples: Configuring Application Identification Properties on page 688](#)

Configuring Automatic Download of Application Package Updates

You can set up automatic downloading of application package updates. To configure downloads, include the **download** statement at the **[edit services application-identification]** hierarchy level:

```
download {  
  automatic {  
    interval hour;  
    start-time time;  
  }  
  url url;  
}
```

You can include the following download statements:

- **download**—Define download properties.
- **automatic**—Set **start-time** value and **interval** in hours for automatic downloads. The default **start-time** is **0:00** and the range is from 0:00 through 24:00. The default **interval** is **24** and the range is from 1 through 168.
- **url**—Specify the download URL.

Related Documentation

- [APPID Overview on page 669](#)
- [Defining an Application Identification on page 674](#)
- [Examples: Configuring Application Identification Properties on page 688](#)

Tracing APPID Operations

Tracing operations track all adaptive services operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

By default, no events are traced. If you include the **traceoptions** statement at the **[edit services application-identification]** hierarchy level, the default tracing behavior is as follows:

- Important events are logged in a file called **serviced** located in the **/var/log** directory.
- When the file **serviced** reaches 128 kilobytes (KB), it is renamed **serviced.0**, then **serviced.1**, and so on, until there are three trace files. Then the oldest trace file (**serviced.2**) is overwritten. (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)
- Only the user who configures the tracing operation can access the log files.
- To display the end of the log, issue the **show log serviced | last** operational mode command:

```
[edit]  
user@host# run show log serviced | last
```


You cannot change the directory (**/var/log**) in which trace files are located. However, you can customize the other trace file settings by including the following statements:

```
file filename <files number> <match regex> <size size> <(world-readable |
no-world-readable>;
flag {
  all;
}
```

You configure these statements at the **[edit services application-identification traceoptions]** hierarchy level.

These statements are described in the following sections:

- [Configuring the APPID Log Filename on page 687](#)
- [Configuring the Number and Size of APPID Log Files on page 687](#)
- [Configuring Access to the Log File on page 687](#)
- [Configuring a Regular Expression for Lines to Be Logged on page 688](#)
- [Configuring the Tracing Flags on page 688](#)

Configuring the APPID Log Filename

By default, the name of the file that records trace output is **serviced**. You can specify a different name by including the **file** statement at the **[edit services application-identification traceoptions]** hierarchy level:

```
file filename;
```

Configuring the Number and Size of APPID Log Files

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed **filename.0**, then **filename.1**, and so on, until there are three trace files. Then the oldest trace file (**filename.2**) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit services application-identification traceoptions]** hierarchy level:

```
file files number size size;
```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (**filename**) reaches 2 MB, **filename** is renamed **filename.0**, and a new file called **filename** is created. When the new **filename** reaches 2 MB, **filename.0** is renamed **filename.1** and **filename** is renamed **filename.0**. This process repeats until there are 20 trace files. Then the oldest file (**filename.19**) is overwritten by the newest file (**filename.0**).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

Configuring Access to the Log File

By default, only the user who configures the tracing operation can access log files.

To specify that any user can read all log files, include the **file world-readable** statement at the **[edit services application-identification traceoptions]** hierarchy level:

```
file world-readable;
```

To explicitly set the default behavior, include the **file no-world-readable** statement at the **[edit services application-identification traceoptions]** hierarchy level:

```
file no-world-readable;
```

Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including the **match** statement at the **[edit services application-identification traceoptions file filename]** hierarchy level and specifying a regular expression (regex) to be matched:

```
file filename match regex;
```

Configuring the Tracing Flags

By default, if the **traceoptions** configuration is present, only important events are logged. You can configure the trace operations to be logged by including the following statements at the **[edit services application-identification traceoptions]** hierarchy level:

```
flag {  
  all;  
}
```

Currently, the only supported flag is **all**, which instructs the router to trace all operations.

Related Documentation

- [APPID Overview on page 669](#)
- [Defining an Application Identification on page 674](#)
- [Examples: Configuring Application Identification Properties on page 688](#)

Examples: Configuring Application Identification Properties

The following examples show an address-based application identification configuration:

```
[edit services application-identification]  
rule rule1 {  
  application-name test2;  
  address 1 {  
    source {  
      ip 10.110.1.1/16;  
      port-range {  
        tcp 1110-1150;  
      }  
    }  
  }  
  destination {  
    ip 10.11.1.1/16;  
    port-range {  
      tcp 111-1100;  
    }  
  }  
}
```

```
    }  
  }  
  order 1;  
}  
}  
}  
[edit services application-identification]  
rule-set rs1 {  
  rule rule1;  
}  
profile pf1 {  
  rule-set rs1;  
}  
[edit services]  
service-set sset1 {  
  application-identification-profile pf1;  
}
```

The following examples show application group configuration:

```
[edit services application-identification]  
application-group junos:peer-to-peer {  
  index 5;  
  application-groups {  
    junos:chat;  
    junos:file-sharing;  
    junos:voip;  
  }  
}  
[edit services application-identification]  
application-group junos:voip {  
  index 14;  
  applications {  
    junos:h225ras;  
    junos:h225sgn;  
    junos:mgcp;  
    junos:sip;  
  }  
}
```

The following examples show application identification for nested application configuration:

```
nested-application nested1 {  
  type nested1;  
  index 65345;  
  protocol HTTP;  
  signature nestedcust001 {  
    member m01 {  
      context http-url-parsed;  
      pattern .*nested.*;  
      direction any;  
    }  
    maximum-transactions 2;  
    order 3825;  
  }
```


Detecting Suspicious and Anomalous Network Traffic Using IDP

- IDP Overview on page 691
- Best-Effort Application Identification of DPI-Serviced Flows on page 693

IDP Overview



NOTE: Starting with Junos OS Release 12.1, all interface-style services are supported for dynamic Point-to-Point Protocol over Ethernet (PPPoE) subscribers on all MX Series routers with modular Modular Port Concentrators (MPCs).

The Junos Application Aware (previously known as Dynamic Application Awareness for the Junos OS) set of services adds support for the intrusion detection and prevention (IDP) functionality using Deep Packet Inspection (DPI) technology to Juniper Networks MX Series 3D Universal Edge Routers equipped with Multiservices Dense Port Concentrators (MS-DPCs) and M120 or M320 Multiservices Edge Routers equipped with Multiservices 400 PICs.

The IDP functionality is already supported on Juniper Networks SRX Series Services Gateways running Junos OS and is described in the *Junos OS Intrusion Detection and Prevention (IDP) Library for Security Devices*. Starting with Junos OS Release 11.3, support for the IDP functionality is extended to T320, T640, and T1600 routers. In addition, multiple IDP detectors are now supported on the M120, M320, and MX Series routers with Enhanced III Flexible PIC Concentrators (FPCs).



NOTE: In the export version of JUNOS, signature download is not expected to work for AppID and IDP features in Junos Application Aware. In order to make it work, you must additionally install the Crypto Software Suite.

The same CLI statements and commands are used on all platforms with the following caveats:

- **Service sets**—IDP is incorporated as a component of service sets only on the specified Juniper Networks T Series, M Series and MX Series routers. IDP depends on application identification services (APPID) for definition and detection of some Layer 7 applications. Before configuring an IDP policy, you must download the APPID application package. Only one service set can be applied to a single interface when the APPID functionality is used.
- **Multiple IDP detectors**—Except for the maximum number of decoder binary instances (4) that are loaded into the process space, multiple IDP detectors on the M120, M320, and MX Series routers function in a similar way to the existing IDP detector support on SRX Series devices. To view the current policy and the corresponding detector version, use the **show security idp status detail** command.

To configure IDP properties, include statements at the **[edit security idp]** hierarchy level. In general, you configure IDP processes by including the **idp-policy** statement at the **[edit system processes]** hierarchy level. For use in T Series, M Series and MX Series applications, you then reference this configuration by including the **idp-profile** statement at the **[edit services service-set]** hierarchy level. To configure SNMP IDP objects, include the **idp** statement at the **[edit snmp health-monitor]** hierarchy level. The operational commands for monitoring and regulating IDP activity are the **clear security idp**, **request security idp**, and **show security idp** commands.

To configure the source IP address for downloading security packages, use the command **set security idp security-package source-address ip-address** because it is not possible to download security packages if the router uses private addressing on its outgoing interface. The source address should be a valid IP address on the node.



NOTE: On T Series, M Series and MX Series routers, the IDP **ip-action** statement is supported on TCP, UDP, and ICMP flows. When the **ip-action target** is **service**, the **ip-action** flow is applied if the traffic matches the values specified for the source port, destination port, source address, and destination address. However, for ICMP flows, the destination port is 0, so that any ICMP flow matching the source port, source address, and destination address would be blocked. For more information about the **ip-action** statement, see the *Junos OS CLI Reference*.

When the Multiservices PIC configured for a service set is either administratively taken offline or undergoes a failure, all the traffic entering the configured interface with an IDP service set would be dropped without notification. To avoid this traffic loss, include the **bypass-traffic-on-pic-failure** statement at the **[edit services service-set service-set-name service-set-options]** hierarchy level. When this statement is configured, the affected packets are forwarded in the event of a Multiservices PIC failure or offlining, as though interface-style services were not configured.



NOTE: Data channel applications for protocols such as FTP, TFTP, RTSP, and SIP are not in the same application group as their control channel applications. For example, control channel application `junos:ftp` is in the group `junos:file-server` but the corresponding data application `junos:system:ftp-data` is not in any group.



NOTE: Because the extension-provider package framework lacks aggressive constraint checks, you should not set the `policy-db-size` statement at the `[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]` hierarchy level to a high value. For Junos Application Aware configurations, the recommended values for the extension-provider package options at this hierarchy level are as follows:

- `control-cores = 1`
- `data-cores = 7`
- `object-cache-size = 1280` (for Multiservices 400 PIC and Multiservices DPC)
- `policy-db-size = 200`
- Include these package values: `jservices-idp`, `jservices-appid`, `jservices-llpdf`, `jservices-aacl`

For more information about this configuration, see the following topics in the *SDK Applications Configuration Guide and Command Reference*:

- *Configuring Control and Data Cores*
- *Configuring Memory Settings*
- *Configuring Packages on the PIC*

Related Documentation

- *Configuring Multiple IDP Detectors*

Best-Effort Application Identification of DPI-Serviced Flows

This topic describes the following information:

- [Features that Support Application-Level Filtering on page 693](#)
- [Best-Effort Application Determination on page 694](#)
- [APPID, AAACL, and L-PDF Processing in Preconvergence Scenarios on page 694](#)

Features that Support Application-Level Filtering

On MX Series routers equipped with Multiservices DPCs and M120 or M320 routers equipped with Multiservices 400 PICs, Intrusion Detection and Prevention (IDP) is accomplished by Deep Packet Inspection (DPI) of TCP, UDP, and ICMP flows. The

application identification (APPID) feature defines applications as members of application groups in TCP/UDP/ICMP traffic. IDP depends on APPID for identification and detection of some Layer 7 applications.

The application-aware access list (AACL) service uses application names and groups as matching criteria for filtering traffic. The service defines the applications and application groups for which statistics are collected for a specific user or interface.

The local policy decision function (L-PDF) enables you to configure properties for statistics output. L-PDF supports a process that regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

Best-Effort Application Determination

Typically, APPID conclusively determines the Layer 7 application associated with a given DPI-serviced flow. In these cases, the application identification is final. Occasionally, APPID is only able to make an initial, inconclusive determination of the Layer 7 application associated with a given flow. This is referred to as a "best-effort" application identification. In such cases, the APPID process continues processing packets on that flow and might subsequently make a conclusive determination of the application associated with that flow. In some cases of best-effort application identification, the flow ends before a final application determination can be made.

APPID, AACL, and L-PDF Processing in Preconvergence Scenarios

The following sections describe APPID, AACL, and L-PDF processing in various stages of application identification for a DPI-serviced flow of TCP/UDP/ICMP traffic.

- [Prior to a Final or Best-Effort Application Identification on page 694](#)
- [Upon Best-Effort Application Identification on page 695](#)
- [While Application Identification Is on a Best-Effort Basis on page 695](#)
- [If a Flow Ends Before an Application Identification Is Made on page 695](#)
- [If a Flow Ends While Application Identification on a Best-Effort Basis on page 695](#)

Prior to a Final or Best-Effort Application Identification

During the time that APPID has not yet made either a final or best-effort determination of the application associated with a given flow, the flow does not contribute to any per-subscriber or per-application statistics collection.

The output of the following operational mode commands includes flows for which APPID has not yet made either a final or best-effort determination of the associated application:

- **show services local-policy-decision-function flows (interface *interface-name* | subscriber *subscriber-name*)**
- **show services application-aware-access-list flows (interface *interface-name* | subscriber *subscriber-name*)**

In the command output, the **Action** field displays "accept" and the **Application** or **Application group** field displays "unknown" for a flow for which APPID has not yet made either a final or best-effort determination of the associated application.

Upon Best-Effort Application Identification

When a best-effort application determination is made, ACL does not apply any ACL term actions configured for that flow. There are a number of reasons for this, one being that the action itself (such as "discard") could make a final application determination impossible. Instead, ACL or L-PDF tracks the flow and accepts all packets for that flow until a final determination is made, at which time the normal ACL or L-PDF actions are fully applied to the flow.

While Application Identification Is on a Best-Effort Basis

During the time that APPID identification of the application associated with a given flow is on a best-effort basis, the flow does not contribute to any per-subscriber or per-application statistics collection.

The output of the following operational mode commands includes flows for which APPID has only made a best-effort determination of the associated application:

- **show services local-policy-decision-function flows** (interface *interface-name* | subscriber *subscriber-name*)
- **show services application-aware-access-list flows** (interface *interface-name* | subscriber *subscriber-name*)

In the command output, the **Action** field displays "accept" and the **Application** or **Application group** field displays "unknown" for a flow for which APPID has only made a best-effort determination of the associated application.

If a Flow Ends Before an Application Identification Is Made

If a flow ends before APPID has made either a final or a best-effort application identification, ACL or L-PDF uses the "unknown" application ID as a final determination and performs any necessary collection, aggregation, and reporting of statistics based on that Layer 7 application. In particular, if the **count** ACL term action is configured for the "application-group-any" application, then the statistics for that flow will be collected and aggregated against the count bucket type, and reported as such.

If a Flow Ends While Application Identification on a Best-Effort Basis

If a flow ends while the application identification is on a best-effort basis, ACL or L-PDF uses that best-effort determination as a final determination. ACL or L-PDF performs any necessary collection, aggregation, and reporting of statistics based on that Layer 7 application. In particular, if the **count** ACL term action is configured for that Layer 7 application, then the statistics for the flow will be collected and aggregated against the ACL or L-PDF statistics. However, in the case of nested applications, ACL and L-PDF will not consider the best-effort determination as final and the nested application will be reported as an unknown application.

**Related
Documentation**

- [Configuring AACL Rules on page 661](#)
- [Configuring Statistics Profiles on page 702](#)
- [aACL-fields on page 1539](#)
- [aACL-statistics-profile on page 1540](#)
- [rule on page 1575](#)
- [services on page 1582](#)
- [term on page 1588](#)
- [then on page 1589](#)

Collecting Statistics and Tracking Data Using L-PDF

- [L-PDF Overview on page 697](#)
- [Best-Effort Application Identification of DPI-Serviced Flows on page 699](#)
- [Configuring Statistics Profiles on page 702](#)
- [Applying L-PDF Profiles to Service Sets on page 705](#)
- [Tracing L-PDF Operations on page 707](#)

L-PDF Overview



NOTE: Starting with Junos OS Release 12.1, all interface-style services are supported for dynamic Point-to-Point Protocol over Ethernet (PPPoE) subscribers on all MX Series routers with modular Modular Port Concentrators (MPCs).

Starting with Junos OS Release 12.1, the local policy decision function (L-PDF) plug-in can offload flows to the Packet Forwarding Engine. Offloading is supported only on MX Series routers with Modular Port Concentrators (MPCs) and accomplished using the Juniper Forwarding Mechanism (JFM). JFM allows services flows to be offloaded to the Packet Forwarding Engine. However, 5-tuple flows cannot be offloaded. Apart from the local L-PDF plug-in, offloading is supported on the packet-triggered subscribers and policy control (PTSP) plug-in. The `show services application-aware-access-list flows subscriber subscriber-name` command displays offload status.

Local policy decision functionality for application-related services adds support for a new process that regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces. This functionality is collectively named the local policy decision function (L-PDF). L-PDF is supported on:

- MX Series routers equipped with Multiservices DPCs.
- M120 or M320 routers equipped with Multiservices 400 PICs.

- Aggregated Multiservices (AMS) PICs.

Multiple `ms-` interfaces can be bundled together in an AMS PIC interface, which causes the traffic destined for this AMS group to be distributed over the member services PICs of the group. Junos OS Trio chipsets enable the calculation of a symmetric hash for the forward and reverse flows, and support a microcode map in the forwarding plane. This capability enables load-balancing of traffic across various services PICs in an AMS group. Starting with Junos OS Release 12.1, `ams-` interfaces enable an N:1 redundancy mechanism to cluster together N number of **`ms-` interfaces** in an AMS group that supports load sharing.

Starting with Junos OS Release 11.3, local L-PDF that resides on the services PIC is supported on T320, T640, and T1600 routers. The application identification (APPID) service defines the applications and how they are grouped. The application-aware access list (AACL) service defines the applications and application groups for which statistics are collected for a specific user or interface. The L-PDF configuration defines the way in which the statistics are output.

To configure properties for statistics output, include the **`policy-decision-statistics-profile`** statement at the **`[edit accounting-options]`** hierarchy level. A new **`traceoptions`** configuration is available at the **`[edit system services local-policy-decision-function]`** hierarchy level. To configure a dynamic profile to attach a specified service set to an interface, include the **`service`** statement at the **`[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family inet]`** hierarchy level. To attach a service set to a static interface, include the **`service-set service-set-name`** statement at the **`[edit interfaces interface-name unit logical-unit-number family inet service (input | output)]`** hierarchy level. For more information on service sets, see *Service Set Properties*.

The following related operational commands are supported:

- **`show services local-policy-decision-function flows`**
- **`show/clear services local-policy-decision-function statistics`**
- **`show/clear services application-aware-access-list statistics`**

For more information on the CLI configuration, see the *Local Policy Decision Function*. For more information on the operational commands, see the [CLI Explorer](#).



NOTE: Because the Junos OS extension-provider package (variously known as JSF, MP-SDK, and eJunos in releases earlier than 12.3) lacks aggressive constraint checks, you should not set the `policy-db-size` statement at the [edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider] hierarchy level to a high value. For Junos Application Aware (previously known as Dynamic Application Awareness) configurations, the recommended values for the extension-provider package options at this hierarchy level are as follows:

- `control-cores = 1`
- `data-cores = 7`
- `object-cache-size = 1280` (for Multiservices 400 PIC and Multiservices DPC)
- `policy-db-size = 200`
- Include these package values: `jservices-idp`, `jservices-appid`, `jservices-llpdf`, `jservices-aacl`

For more information about this configuration, see the following topics in the *SDK Applications Configuration Guide and Command Reference*:

- *Configuring Control and Data Cores*
- *Configuring Memory Settings*
- *Configuring Packages on the PIC*

Related Documentation

- [Best-Effort Application Identification of DPI-Serviced Flows on page 658](#)
- [Configuring Statistics Profiles on page 702](#)
- [Applying L-PDF Profiles to Service Sets on page 705](#)
- [Tracing L-PDF Operations on page 707](#)

Best-Effort Application Identification of DPI-Serviced Flows

This topic describes the following information:

- [Features that Support Application-Level Filtering on page 699](#)
- [Best-Effort Application Determination on page 700](#)
- [APPID, AACL, and L-PDF Processing in Preconvergence Scenarios on page 700](#)

Features that Support Application-Level Filtering

On MX Series routers equipped with Multiservices DPCs and M120 or M320 routers equipped with Multiservices 400 PICs, Intrusion Detection and Prevention (IDP) is accomplished by Deep Packet Inspection (DPI) of TCP, UDP, and ICMP flows. The application identification (APPID) feature defines applications as members of application

groups in TCP/UDP/ICMP traffic. IDP depends on APPID for identification and detection of some Layer 7 applications.

The application-aware access list (AACL) service uses application names and groups as matching criteria for filtering traffic. The service defines the applications and application groups for which statistics are collected for a specific user or interface.

The local policy decision function (L-PDF) enables you to configure properties for statistics output. L-PDF supports a process that regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

Best-Effort Application Determination

Typically, APPID conclusively determines the Layer 7 application associated with a given DPI-serviced flow. In these cases, the application identification is final. Occasionally, APPID is only able to make an initial, inconclusive determination of the Layer 7 application associated with a given flow. This is referred to as a "best-effort" application identification. In such cases, the APPID process continues processing packets on that flow and might subsequently make a conclusive determination of the application associated with that flow. In some cases of best-effort application identification, the flow ends before a final application determination can be made.

APPID, AACL, and L-PDF Processing in Preconvergence Scenarios

The following sections describe APPID, AACL, and L-PDF processing in various stages of application identification for a DPI-serviced flow of TCP/UDP/ICMP traffic.

- [Prior to a Final or Best-Effort Application Identification on page 700](#)
- [Upon Best-Effort Application Identification on page 701](#)
- [While Application Identification Is on a Best-Effort Basis on page 701](#)
- [If a Flow Ends Before an Application Identification Is Made on page 701](#)
- [If a Flow Ends While Application Identification on a Best-Effort Basis on page 701](#)

Prior to a Final or Best-Effort Application Identification

During the time that APPID has not yet made either a final or best-effort determination of the application associated with a given flow, the flow does not contribute to any per-subscriber or per-application statistics collection.

The output of the following operational mode commands includes flows for which APPID has not yet made either a final or best-effort determination of the associated application:

- **show services local-policy-decision-function flows (interface *interface-name* | subscriber *subscriber-name*)**
- **show services application-aware-access-list flows (interface *interface-name* | subscriber *subscriber-name*)**

In the command output, the **Action** field displays "accept" and the **Application** or **Application group** field displays "unknown" for a flow for which APPID has not yet made either a final or best-effort determination of the associated application.

Upon Best-Effort Application Identification

When a best-effort application determination is made, ACL does not apply any ACL term actions configured for that flow. There are a number of reasons for this, one being that the action itself (such as "discard") could make a final application determination impossible. Instead, ACL or L-PDF tracks the flow and accepts all packets for that flow until a final determination is made, at which time the normal ACL or L-PDF actions are fully applied to the flow.

While Application Identification Is on a Best-Effort Basis

During the time that APPID identification of the application associated with a given flow is on a best-effort basis, the flow does not contribute to any per-subscriber or per-application statistics collection.

The output of the following operational mode commands includes flows for which APPID has only made a best-effort determination of the associated application:

- **show services local-policy-decision-function flows** (interface *interface-name* | subscriber *subscriber-name*)
- **show services application-aware-access-list flows** (interface *interface-name* | subscriber *subscriber-name*)

In the command output, the **Action** field displays "accept" and the **Application** or **Application group** field displays "unknown" for a flow for which APPID has only made a best-effort determination of the associated application.

If a Flow Ends Before an Application Identification Is Made

If a flow ends before APPID has made either a final or a best-effort application identification, ACL or L-PDF uses the "unknown" application ID as a final determination and performs any necessary collection, aggregation, and reporting of statistics based on that Layer 7 application. In particular, if the **count** ACL term action is configured for the "application-group-any" application, then the statistics for that flow will be collected and aggregated against the count bucket type, and reported as such.

If a Flow Ends While Application Identification on a Best-Effort Basis

If a flow ends while the application identification is on a best-effort basis, ACL or L-PDF uses that best-effort determination as a final determination. ACL or L-PDF performs any necessary collection, aggregation, and reporting of statistics based on that Layer 7 application. In particular, if the **count** ACL term action is configured for that Layer 7 application, then the statistics for the flow will be collected and aggregated against the ACL or L-PDF statistics. However, in the case of nested applications, ACL and L-PDF will not consider the best-effort determination as final and the nested application will be reported as an unknown application.

Related Documentation

- [Configuring ACL Rules on page 661](#)
- [Configuring Statistics Profiles on page 702](#)
- [acl-fields on page 1539](#)

- [aacl-statistics-profile on page 1540](#)
- [rule on page 1575](#)
- [services on page 1582](#)
- [term on page 1588](#)
- [then on page 1589](#)

Configuring Statistics Profiles

The local policy decision function (L-PDF) enables you to configure properties for statistics output. To do this, you create a statistics profile, which configures the files to which statistics records are exported and the format that is exported. There are two configurations you can use to specify the profile, as described in the following subsections:

- [Configuring an L-PDF Statistics Profile on page 703](#)
- [Configuring an AACL Statistics Profile on page 704](#)



NOTE: You must use the same configuration stanza for specifying the profile and the file selection. If configurations are committed in both hierarchies, the one at the `[edit system services local-policy-decision-function]` hierarchy level takes precedence.



NOTE:

- When a session closes before APPID has identified nested applications, the session is treated as a best-effort session and L-PDF does not get the nested application information. In such cases, nested applications will be reported as unknown applications.
- During the time that the application identification (APPID) feature has not yet made a final determination of the application associated with a given flow, the flow does not contribute to any per-subscriber or per-application statistics collection. For more information, see [“Best-Effort Application Identification of DPI-Serviced Flows” on page 658](#).



NOTE: For rms- interfaces, the statistics received from the active Multiservices PICs in the RMS group are combined with the statistics of the reported ended flows kept on the Routing Engine. The aggregated value is written to the statistics file. In the case of AMS interfaces, all the Multiservices PICs consisting of the AMS group reports statistics independently. These statistics are aggregated on the Routing Engine. The Routing Engine runs an independent timer, which on expiry writes the aggregated entry in the statistics file. This method of collection causes the statistics data in the statistics file to be displayed with a small delay.

Configuring an L-PDF Statistics Profile

You can specify an L-PDF statistics profile by including the following configuration at the **[edit accounting-options]** hierarchy level:

```
[edit accounting-options]
policy-decision-statistics-profile profile-name {
  application-aware-access-list-fields [ field-name ];
  file filename;
  files number;
  size bytes;
}
```



NOTE: This configuration method is not the preferred method for configuring Junos Application Aware (previously known as Dynamic Application Awareness) statistics. It is only maintained for backwards compatibility and may be deprecated in a future software release and does not support the use of IPv6 address and prefix length. The new, preferred configuration is found at the **[edit system services local-policy-decision-function]** hierarchy level, as described in “[Configuring an ACL Statistics Profile](#)” on page 704. We encourage you to migrate to the new configuration method.

You specify a profile name to identify the profile and other properties as needed by including the **policy-decision-statistics-profile** statement. The **acl-fields** statement specifies which statistics to collect in an accounting-data log file. This log file is located on the **/var/log** directory on the router. You specify the log file by including the **file filename** statement. The filename is prefixed by the **acl_statistics_** prefix; for example, if you specify the filename **lpdfd**, the log file will be **/var/log/acl_statistics_lpdfd**.

The **application-aware-access-list-fields** statement supports the following options:

- **address**—IP Address
- **application**—Application name
- **application-group**—Application group name
- **input-bytes**—Number of input bytes
- **input-interface**—Input interface name
- **input-packets**—Number of input packets
- **mask**—Netmask
- **output-bytes**—Number of output bytes
- **output-packets**—Number of output packets
- **subscriber-name**—Subscriber name
- **timestamp**—Timestamp
- **vrf-name**—VPN routing and forwarding (VRF) name

For more information on configuring profiles, see the *Network Management Administration Guide for Routing Devices*.

Configuring an ACL Statistics Profile

You can specify an ACL statistics profile by including the following configuration at the **[edit system services]** hierarchy level:

```
local-policy-decision-function {
  statistics {
    file filename {
      archive-sites [ url ];
      files number;
      size bytes;
      transfer-interval minutes;
    }
    aacl-statistics-profile profile-name {
      aacl-fields [ field-name ];
      file filename;
      report-interval minutes;
      record-mode (interim-active-only | interim-full);
    }
    record-type (delta | interim);
  }
}
```

To specify the file properties, include the **file** statement at the **[edit system services local-policy-decision-function statistics]** hierarchy level with a unique filename:

- The **archive-sites** statement specifies one or more URLs for archiving the files. Archiving can be done by using FTP or SCP.
- The **files** statement specifies the maximum number of files that are maintained at one time.
- The **size** statement specifies the maximum size of each file.
- The **transfer-interval** statement specifies the interval between data transfers in minutes.

You specify a profile name to identify the profile and other properties as needed by including the **aacl-statistics-profile** statement. The **aacl-fields** statement specifies which statistics to collect in an accounting-data log file. This log file is located on the **/var/stats/aacl** directory on the router. You specify the log file by including the **file filename** statement.

The **aacl-fields** statement supports the following options:

- **address**—IP Address
- **all-fields**—All available fields
- **application**—Application name
- **application-group**—Application group name
- **input-bytes**—Number of input bytes
- **input-interface**—Input interface name

- **ipv6-address**—IPv6 address
- **ipv6-prefix-length**—Prefix length associated with the displayed IPv6 address
- **input-packets**—Number of input packets
- **mask**—Netmask
- **output-bytes**—Number of output bytes
- **output-packets**—Number of output packets
- **subscriber-name**—Subscriber name
- **timestamp**—Timestamp
- **vrf-name**—VPN routing and forwarding (VRF) name

The **record-type** statement specifies whether a record is **delta** or **interim**; **delta** is the default setting. The **report-interval** statement specifies the reporting interval in minutes; the default setting is 15 minutes and the range is 5 through 1440 minutes. The **record-mode** statement specifies how the statistics are reported for each reporting interval; the default setting is **interim-full** and reports all available statistics. To report only statistics that have changed for the reporting interval, use the **interim-active-only** setting.



NOTE: The IPv6 fields (**ipv6-address** and **ipv6-prefix-length**) are not supported for **record-type delta**. The IPv6 fields are supported for **record-type interim** only, meaning that the fields are restricted to the S- (Login) record.

For more information on configuring profiles, see the *Network Management Administration Guide for Routing Devices*.

Related Documentation

- [L-PDF Overview on page 697](#)
- [Best-Effort Application Identification of DPI-Serviced Flows on page 658](#)
- [Applying L-PDF Profiles to Service Sets on page 705](#)
- [Tracing L-PDF Operations on page 707](#)

Applying L-PDF Profiles to Service Sets

You can optionally apply policy decision statistics profiles as part of a service-set definition. To do this, you include the **policy-decision-statistics-profile** statement at the **[edit services service-set *service-set-name*]** hierarchy level:

```
policy-decision-statistics-profile profile-name;
```



NOTE: To provide high availability for the policy decision statistics, associate the service-set definition with a redundant services PIC (rsp) interface.

You can include only one profile name in the specification for the **application-aware access-list** statement.

The following example shows a sample configuration for attachment of an L-PDF statistics profile:

```
services {
  service-set test_aacl_sset {
    aacl-rules aacl_rule;
    policy-decision-statistics-profile {
      pdf_stats_prof;
    }
    interface-service {
      service-interface ms-0/3/0.0;
    }
  }
}
```



NOTE: Only one service set can be applied to a single interface when L-PDF functionality is used.

The following example shows a sample configuration for attachment of a service set to a static interface:

```
interfaces {
  fe-0/0/0 {
    vlan-tagging;
    unit 1 {
      vlan-id 1;
      family inet {
        service {
          input {
            service-set test_aacl_sset;
          }
          output {
            service-set test_aacl_sset;
          }
        }
      }
      address 10.1.1.1/24;
    }
  }
}
```



NOTE: The `session-offload` statement at the `[edit chassis fpc slot-number pic number adaptive-services service-package extension-provider]` hierarchy level controls session offload behavior for Multiservices DPCs on MX Series routers. It controls session offload on a per-device basis, where a device is a Multiservices interface (`ms-fpc-pic-port`). Currently, the session offload function is supported for at most one Multiservices interface. When the offload function is enabled, it is strongly recommended that you limit Junos Application Aware (previously known as Dynamic Application Awareness) features to that Multiservices interface.

The default is to not offload any sessions. For more information on chassis configuration, see the *Junos OS Administration Library for Routing Devices*.

Related Documentation

- [L-PDF Overview on page 697](#)
- [Best-Effort Application Identification of DPI-Serviced Flows on page 658](#)
- [Configuring Statistics Profiles on page 702](#)
- [Tracing L-PDF Operations on page 707](#)

Tracing L-PDF Operations

Tracing operations track L-PDF operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

By default, no events are traced. If you include the `traceoptions` statement at the `[edit system services local-policy-decision-function]` hierarchy level, you can customize the trace file settings:

```
traceoptions {
  file filename <files number> <size size>;
  flag flag;
}
```

The flags track the following information:

- **all**—Everything
- **configuration**—Configuration traces
- **database**—Database traces
- **general**—Miscellaneous traces
- **gres**—Graceful Routing Engine switchover (GRES) traces
- **ptsp-statistics**—PTSP statistics traces
- **rtsock**—Routing socket traces
- **statistics**—Statistics traces
- **subscriber**—Subscriber traces

**Related
Documentation**

- [L-PDF Overview on page 697](#)
- [Best-Effort Application Identification of DPI-Serviced Flows on page 658](#)
- [Configuring Statistics Profiles on page 702](#)
- [Applying L-PDF Profiles to Service Sets on page 705](#)

PART 13

Configuring Link and Multilink Services Interfaces

- Overview on page 711
- Achieving Greater Bandwidth, Load Balancing, and Redundancy with Multilink Bundles on page 719
- Configuring the Physical and Logical Interfaces in a Multilink Configuration on page 725
- Bundling Multiple PPP Links on a Single Link Using MLPPP on page 747
- Bundling Multiple Frame Relay DLCIs into a Single Link Using MLFR on page 777
- Configuring Additional Services on Link Services Interfaces on page 801

Overview

- [Link and Multilink Services Overview on page 711](#)
- [Multilink and Link Services PICs Overview on page 714](#)
- [Multilink Interfaces on Channelized MICs Overview on page 715](#)
- [Multilink and Link Services Logical Interface Configuration Overview on page 717](#)

Link and Multilink Services Overview

Multilink-based protocols enable you to split, recombine, and sequence datagrams across multiple logical data links. The goal of a multilink operation is to coordinate multiple independent links between a fixed pair of systems, providing a virtual link with greater bandwidth than any of the members.

The Juniper Networks Junos operating system (Junos OS) supports several multilink-based protocols (such as MLPPP, FRF.15, and FRF.16) on the services PICs such as the Multilink Services PIC, the Link Services PIC, and the link services intelligent queuing (IQ) and voice services configured on the Adaptive Services (AS) and MultiServices PICs. For more information about link services IQ, see [“Layer 2 Service Package Capabilities and Interfaces” on page 543](#). For more information about voice services, see [“Configuring Services Interfaces for Voice Services” on page 622](#).

Starting with Junos OS Release 12.1, the following channelized MICs on MX240, MX480, and MX960 routers support Multilink Point-to-Point Protocol (MLPPP)-based services:

- 4-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-4CHOC3-2CHOC12)
- 8-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-8CHOC3-4CHOC12)
- 8-port Channelized DS3/E3 MIC (MIC-3D-8CHDS3-E3-B)

For more information about Multilink Point-to-Point Protocol (MLPPP)-based services MICs, see [“Multilink Interfaces on Channelized MICs Overview” on page 715](#).



NOTE: The ml- interface type is used to configure interfaces on the Multilink Services PIC and does not support class-of-service (CoS) features. The ls- interface type is used for limited CoS configurations on the Link Services PIC, and the lsq- interface type is used for full CoS configurations on the Adaptive Services and MultiServices PICs. The bundle interfaces are configured on the Multiservices DPC as link services IQ (lsq) interfaces and virtual LSQ redundancy (rlsq) interfaces.

For link services IQ (lsq) interfaces, Junos OS CoS components are fully supported and are handled normally on M Series and T Series routers, as described in the *Class of Service Feature Guide for Routing Devices*. For more information on link services IQ configuration, see [“Layer 2 Service Package Capabilities and Interfaces” on page 543](#).

The Link Services and Multilink Services PICs support the following encapsulation types:

- Multilink Point-to-Point Protocol (MLPPP)
- Multilink Frame Relay (MLFR)

Starting with Junos OS Release 12.1, support for the following encapsulation types and protocols has been extended to the MX240, MX480, and MX960 routers with Multiservices DPCs:

- Multilink Point-to-Point Protocol (MLPPP)
- Multiclass MLPPP
- Multilink Frame Relay (MLFR) end-to-end (FRF.15)
- Multilink Frame Relay (MLFR) UNI NNI (FRF.16) (also referred to as MFR)
- Compressed Real-Time Transport Protocol (CRTP)

MLPPP enables you to bundle multiple PPP links into a single logical link. MLFR enables you to bundle multiple Frame Relay data-link connection identifiers (DLCIs) into a single logical link. MLPPP and MLFR provide service option granularity between low-speed T1 and E1 services and higher-speed T3 and E3 services. You use MLPPP and MLFR to increase bandwidth in smaller, more cost-effective increments. In addition to providing incremental bandwidth, bundling multiple links can add a level of fault tolerance to your dedicated access service, because you can implement bundling across multiple PICs, protecting against the failure of any single PIC.



NOTE: Even if the PIC can support up to 4xDS3 total throughput, each aggregate can only run a volume of traffic equal to one DS3 in bandwidth. Aggregating DS3 links is not supported.

At the logical unit level, the Multilink Services and Link Services PICs support the MLPPP and MLFR Frame Relay Forum (FRF) 15 encapsulation types. At the physical interface level, the Link Services PIC also supports the MLFR FRF.16 encapsulation type.

MLPPP and MLFR FRF.15 are supported on interface types **ml-fpc/pic/port**, **ls-fpc/pic/port**, and **lsq-fpc/pic/port**. For MLFR FRF.15, multiple permanent virtual circuits (PVCs) are combined into one aggregated virtual circuit (AVC). This provides fragmentation over multiple PVCs on one end and reassembly of the AVC on the other end.

MLFR FRF.16 is supported on a channelized interface, **ls-fpc/pic/port:channel**, which denotes a single MLFR FRF.16 bundle. For MLFR FRF.16, multiple links are combined to form one logical link. Packet fragmentation and reassembly occur on a per-VC basis. Each bundle can support multiple VCs. Link Services PICs can support up to 256 DLCIs per MLFR FRF.16 bundle. The physical connections must be E1, T1, channelized DS3-to-DS1, channelized DS3-to-DS0, channelized E1, channelized STM1, or channelized IQ interfaces. When you bundle channelized interfaces using the link services interface, the channelized interfaces require M Series Enhanced Flexible PIC Concentrators (FPCs).



NOTE: When running MLPPP or MLFR on a non-QPP interface, you cannot mix logical units that are members of an aggregate with logical units configured using other families, such as **inet**. For example, the following configuration is not valid:

```
interface e3-0/0/0 {
  encapsulation frame-relay;
  unit 99 {
    dlci 99;
    family mlfr-end-to-end {
      bundle ls-0/0/0.1;
    }
  }
  unit 100 { ## mixes mlfr with family inet
    dlci 100;
    family inet {
      address 192.168.164.53/30;
    }
  }
}
```

The standards for MLPPP, MLFR FRF.15, and MLFR FRF.16 are defined in the following specifications:

- RFC 1990, *The PPP Multilink Protocol (MP)*
- FRF.15, *End-to-End Multilink Frame Relay Implementation Agreement*
- FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*



NOTE: Endpoint Discriminator Class compatibility checking is enabled on MLPPP interfaces. Prior to Junos OS Release 8.0, when a Juniper Networks router received an unsupported Endpoint Discriminator Class message from an MLPPP session peer, it returned an ACK response.

Related Documentation

- [Understanding MLPPP Bundles and Link Fragmentation and Interleaving \(LFI\) on Serial Links on page 719](#)
- [Example: Configuring an MLPPP Bundle on page 750](#)
- [Example: Configuring Multilink Frame Relay FRF.15 on page 783](#)
- [Example: Configuring Multilink Frame Relay FRF.16 on page 779](#)

Multilink and Link Services PICs Overview

Each Multilink Services or Link Services PIC can support a number of *bundles*. A bundle can contain up to eight individual *links*.

For Multilink Services PICs, the links can be T1, E1, or DS0 physical interfaces, and each link is associated with a logical unit number that you configure. For Link Services PICs, the links can be E1, T1, channelized DS3-to-DS1, channelized DS3-to-DS0, channelized E1, channelized STM1 interfaces, or channelized IQ interfaces. For MLFR FRF.16 bundles, each link is associated with a channel number that you configure.

You must configure a link before it can join a bundle. Each bundle should consist solely of one type of link; the mixing of physical interfaces of differing speeds within a bundle is not supported.



NOTE: On M Series Multiservice Edge Routers, only one DS3 link is allowed in an MLFR bundle. MLPPP bundles can include two DS3 links.

Three versions of Multilink Services and three versions of Link Services PICs are available, as shown in [Table 27 on page 714](#). The PIC hardware is identical, except for different faceplates that enable you to identify which version you are installing. The software limits the unit numbers and maximum number of physical interfaces you assign to the PIC.

Table 27: Multilink and Link Services PIC Capacities

PIC Capacity	Unit Numbers	Maximum Number of T1/DS0 Interfaces	Maximum Number of E1 Interfaces
4-bundle PIC	0 through 3	32 links	32 links
32-bundle PIC	0 through 31	256 links	219 links
128-bundle PIC	0 through 127	292 links	219 links

A single PIC can support an aggregate bandwidth of 450 megabits per second (Mbps).

You can configure a larger number of links, but the Multilink Services and Link Services PICs can reliably process only 450 Mbps of traffic. A higher rate of traffic might degrade performance.



NOTE: In Junos OS releases 9.0 and above you are not allowed to configure a unit number greater than the maximum unit number available on your link services PIC. Attempting to do so will cause an error message.

Related Documentation

- [Link and Multilink Services Overview on page 711](#)
- [Multilink Interfaces on Channelized MICs Overview on page 715](#)
- [Multilink and Link Services Logical Interface Configuration Overview on page 717](#)
- [Understanding MLPPP Bundles and Link Fragmentation and Interleaving \(LFI\) on Serial Links on page 719](#)
- [Configuring the Number of Bundles on Link Services PICs on page 720](#)
- [Configuring the Links in a Multilink or Link Services Bundle on page 721](#)

Multilink Interfaces on Channelized MICs Overview

Multiservices Modular Interface Cards (MICs) enable you to perform multiple services on the same MIC by configuring a set of services and applications such as voice services and Layer 2 Tunneling Protocol (L2TP) services. On Juniper Networks MX Series 3D Universal Edge Routers, the Multiservices DPC provides essentially the same capabilities as the Multiservices PIC. The interfaces on both platforms are configured in the same way. The Multilink interfaces are hosted on a channelized MIC. The bundle interfaces are configured on Multiservices DPC as virtual LSQ redundancy (rlsq) interfaces.

Starting with Junos OS Release 12.1, the following channelized MICs on MX240, MX480, and MX960 routers support Multilink Point-to-Point Protocol (MLPPP)-based services:

- 4-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-4CHOC3-2CHOC12)
- 8-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-8CHOC3-4CHOC12)
- 8-port Channelized DS3/E3 MIC (MIC-3D-8CHDS3-E3-B)

The following encapsulations, interfaces, protocol, and packet types are supported on the aforementioned MICs:

- Multilink Point-to-Point Protocol (MLPPP)—Supports Priority-based Flow Control (PFC) for data packets and Link Control Protocol (LCP) for control packets. Compressed Real-Time Transport Protocol (CRTP) and Multiclass MLPPP are supported for both data and control packets.
- Multilink Frame Relay (MLFR) end-to-end (FRF.15)—Supports Ethernet Local Management Interface (LMI), Consortium LMI (C-LMI), and Link Integrity Protocol (LIP) for data and control packets.

- Multilink Frame Relay (MFR) UNI NNI (FRF.16)—Supports Ethernet Local Management Interface (LMI), Consortium LMI (C-LMI), and Link Integrity Protocol (LIP) for data and control packets.
- Link fragmentation and interleaving (LFI) non multilink MLPPP and MLFR packets.

Layer 2 services and voice services functionality are implemented on the Multiservices Dense Port Concentrators which supports the following two kinds of traffic that are routed by the Packet Forwarding Engine:

- Customer-end to provider-end (also, known as customer traffic)—Here, the Multilink fragments from the customer end arrive at the Multiservices interfaces configured on the channelized MIC. These fragments are then transmitted to the Multiservices DPC for Layer 2 processing such as CoS and are reassembled by the Multiservices software running on the Multiservices DPC. These reassembled packets are sent to the Packet Forwarding Engine where they go through the regular router lookup process and are finally sent over the Internet to the provider end. The voice packets also go through the same process.
- Provider-end to customer-end (also, known as Internet traffic)—Here, the data packets that are sent from the Internet provider end are received at any generic ingress interface in the Packet Forwarding Engine. These packets are then sent to the Multiservices DPC for Layer 2 processing. The Multiservices software running on Multiservices DPC fragment these data packets and send it to the Packet Forwarding Engine. These Multilink fragments are sent over the channelized MIC interfaces to the customer end. The voice packets also go through the same process.



NOTE: All the features that are supported on Multilink and Link Services PICs are also supported on the Multilink Services or Link Services MICs. For more information about Multilink and Link Services PICs, see [“Multilink and Link Services PICs Overview” on page 714.](#)

Support for the following encapsulations, interfaces, protocol, and packet types are now extended to the aforementioned MICs:

- Multilink Point-to-Point Protocol (MLPPP)—Supports priority-based flow control (PFC) for data packets and Link Control Protocol (LCP) for control packets. Compressed Real-Time Transport Protocol (CRTP) and multiclass MLPPP are supported for both data and control packets.
- Multilink Frame Relay (MLFR) end-to-end (FRF.15)—Supports Ethernet Local Management Interface (LMI) and Consortium LMI (C-LMI) for data and control packets.
- Multilink Frame Relay (MLFR) UNI NNI (FRF.16)—Supports Ethernet Local Management Interface (LMI), Consortium LMI (C-LMI), and Link Integrity Protocol (LIP) for data and control packets.
- Link fragmentation and interleaving (LFI) on multilink MLPPP and MLFR packets—Reduces delay and jitter on links by breaking up large data packets and interleaving delay-sensitive voice packets with the resulting smaller packets.

- Related Documentation**
- [Link and Multilink Services Overview on page 711](#)
 - [Example: Configuring Link Interfaces on Channelized MICs on page 735](#)
 - [Example: Configuring an MLPPP Bundle on page 750](#)
 - [Example: Configuring Multilink Frame Relay FRF.15 on page 783](#)
 - [Example: Configuring Multilink Frame Relay FRF.16 on page 779](#)

Multilink and Link Services Logical Interface Configuration Overview

You configure multilink and link services interface properties at the logical unit level. Default settings for multilink and link services logical interface properties are described in [“Default Settings for Multilink and Link Services Logical Interfaces” on page 717](#).

For general information about logical unit properties or **family inet** properties, see the *Junos OS Network Interfaces Library for Routing Devices*. For information about multilink and link services properties you configure at the **family inet** hierarchy level, see [“Configuring the Links in a Multilink or Link Services Bundle” on page 721](#).



NOTE: On DS0, E1, or T1 interfaces in LSQ bundles, you can configure the **bandwidth** statement, but the router does not use the bandwidth value if the interfaces are included in an MLPPP or MLFR bundle. The bandwidth is calculated internally according to the time slots, framing, and byte-encoding of the interface. For more information about logical interface properties, see the *Junos OS Network Interfaces Library for Routing Devices*.

Default Settings for Multilink and Link Services Logical Interfaces

Table 28 on page 717 lists the default settings for multilink and link services statements, together with the other permitted values or value ranges.

Table 28: Multilink and Link Services Logical Interface Statements

Option	Default Value	Possible Values
DLCI	None	16 through 1022
Drop timeout period	500 ms for bundles greater than or equal to the T1 bandwidth value and 1500 ms for other bundles.	0 through 2000 milliseconds
Encapsulation	For multilink interfaces, multilink-ppp . For link services interfaces, multilink-frame-relay-end-to-end .	multilink-frame-relay-end-to-end , multilink-ppp
Fragmentation threshold	0 bytes	128 through 16,320 bytes (N×64)

Table 28: Multilink and Link Services Logical Interface Statements (*continued*)

Option	Default Value	Possible Values
Interleave fragments	disabled	enabled, disabled
Minimum links	1 link	1 through 8 links
Maximum received reconstructed unit (MRRU)	1504 bytes	1500 through 4500 bytes
Sequence ID format for MLPPP	24 bits	12 or 24 bits
Sequence ID format for MLFR FRF.15 and FRF.16	12 bits	12 bits

See “[Default Settings for Link Services Interfaces](#)” on [page 726](#) for statements that apply to link services physical interfaces only.

Related Documentation

- [Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 729](#)
- [Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces on page 730](#)
- [Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 732](#)
- [Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces on page 733](#)
- [Configuring MRRU on Multilink and Link Services Logical Interfaces on page 734](#)
- [Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces on page 735](#)
- [Configuring DLCIs on Link Services Logical Interfaces on page 777](#)
- [Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces on page 773](#)

CHAPTER 50

Achieving Greater Bandwidth, Load Balancing, and Redundancy with Multilink Bundles

- [Understanding MLPPP Bundles and Link Fragmentation and Interleaving \(LFI\) on Serial Links on page 719](#)
- [Configuring the Number of Bundles on Link Services PICs on page 720](#)
- [Configuring the Links in a Multilink or Link Services Bundle on page 721](#)
- [Example: Configuring a Link Services Interface with Two Links on page 722](#)

Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links

MX240, MX480, and MX960 3D Universal Edge Routers support MLPPP and MLFR multilink encapsulations. MLPPP enables you to bundle multiple PPP links into a single multilink bundle, and MLFR enables you to bundle multiple Frame Relay data-link connection identifiers (DLCIs) into a single multilink bundle. Multilink bundles provide additional bandwidth, load balancing, and redundancy by aggregating low-speed links, such as T1, E1, and serial links.

You configure multilink bundles as logical units or channels on the link services interface **lsq-0/0/0**:

- With MLPPP and MLFR FRF.15, multilink bundles are configured as logical units on **lsq-0/0/0**—for example, **lsq-0/0/0.0** and **lsq-0/0/0.1**.
- With MLFR FRF.16, multilink bundles are configured as channels on **lsq-0/0/0**—for example, **lsq-0/0/0:0** and **lsq-0/0/0:1**.

After creating multilink bundles, you add constituent links to the bundle. The constituent links are the low-speed physical links that are to be aggregated. You can create 64 multilink bundles, and on each multilink bundle you can add up to 8 constituent links. The following rules apply when you add constituent links to a multilink bundle:

- On each multilink bundle, add only interfaces of the same type. For example, you can add either T1 or E1, but not both.

- Only interfaces with a PPP encapsulation can be added to an MLPPP bundle, and only interfaces with a Frame Relay encapsulation can be added to an MLFR bundle.
- If an interface is a member of an existing bundle and you add it to a new bundle, the interface is automatically deleted from the existing bundle and added to the new bundle.

Configuring a multilink bundle on the two serial links increases the bandwidth by 70 percent from approximately 1 Mbps to 1.7 Mbps and prepends each packet with a multilink header as specified in the FRF.12 standard. To increase the bandwidth further, you can add up to eight serial links to the bundle. In addition to a higher bandwidth, configuring the multilink bundle provides load balancing and redundancy. If one of the serial links fails, traffic continues to be transmitted on the other links without any interruption. In contrast, independent links require routing policies for load balancing and redundancy. Independent links also require IP addresses for each link as opposed to one IP address for the bundle. In the routing table, the multilink bundle is represented as a single interface.

Starting with Junos OS Release 13.3, if you attempt to delete or deactivate a static inline service (**si**) MLPPP bundle interface that is still referenced by a member link interface, which could be PPPoE (**pp0**) or **si** logical interfaces, and commit the configuration, the commit operation fails. You must reactivate such MLPPP bundle interface before committing the settings. Alternatively, you must ensure that member links do not refer a static MLPPP bundle before you delete or deactivate the bundle. This method of deactivation and reactivation of an MLPPP bundle is not applicable for interfaces other than **si**- interfaces, such as link services IQ (**lsq**-) and virtual LSQ redundancy (**rlsq**-) interfaces.

Related Documentation

- [Link and Multilink Services Overview on page 711](#)
- [Multilink Interfaces on Channelized MICs Overview on page 715](#)
- [Example: Configuring Link Interfaces on Channelized MICs on page 735](#)
- [Example: Configuring an MLPPP Bundle on page 750](#)
- [Example: Configuring Multilink Frame Relay FRF.15 on page 783](#)
- [Example: Configuring Multilink Frame Relay FRF.16 on page 779](#)

Configuring the Number of Bundles on Link Services PICs

You can combine MLFR FRF.16, MLPPP, and MLFR FRF.15 bundles on a single Link Services PIC. For a sample configuration, see “[Example: Configuring a Link Services Interface with Two Links](#)” on page 722.

To configure the number of bundles on a Link Services PIC, include the **mlfr-uni-nni-bundles** statement at the **[edit chassis fpc slot-number pic pic-number]** hierarchy level:

```
mlfr-uni-nni-bundles number;
```

Each Link Services PIC can accommodate a maximum of 256 MLFR UNI NNI bundles. For more information, see the *Junos OS Administration Library for Routing Devices*.

A link can associate with one link services bundle only. All Link Services PICs support up to 256 single-link bundles and up to 256 DLCIs. For an example configuration, see the configuration examples.



NOTE: When one or more links in a bundle are put in loopback, reassembly buffering and hence processing are reduced so as to not affect other bundles. This prevents packet loss on other bundles, while reducing the reassembly buffers available for the bundle with looped links.


Related Documentation

- [Example: Configuring a Link Services Interface with Two Links on page 722](#)
- [Example: Configuring a Link Services Interface with MLPPP on page 753](#)
- [Example: Configuring a Link Services Interface with MLFR FRF.15 on page 786](#)
- [Example: Configuring a Link Services PIC with MLFR FRF.16 on page 787](#)
- [Example: Configuring Link and Voice Services Interfaces with a Combination of Bundle Types on page 806](#)

Configuring the Links in a Multilink or Link Services Bundle

To complete a multilink or link services interface configuration, you need to configure both the physical interface and the multilink or link services bundle. For multilink interfaces, you configure the link bundle on the logical unit. For link services interfaces, you configure the link bundle as a channel (see [Figure 26 on page 721](#)). The physical interface is usually connected to networks capable of supporting MLPPP or MLFR (FRF.15 or FRF.16).

Figure 26: Multilink Interface Configuration



The following sample configuration refers to the topology in [Figure 26 on page 721](#) and configures a multilink or link services bundle over a T1 connection (for which the T1 physical interface is already configured).

1. To configure a physical T1 link for MLPPP, include the following statements at the **[edit interfaces t1-fpc/pic/port]** hierarchy level:

```
unit 0 {
  family mlppp {
    bundle (ml-fpc/pic/port | ls-fpc/pic/port);
  }
}
```

You do not need to configure an IP address on this link.

To configure a physical T1 link for MLFR FRF.16, include the following statements at the `[edit interfaces t1-fpc/pic/port]` hierarchy level:

```
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
  family mlfr-uni-nni {
    bundle ls-fpc/pic/port:channel;
  }
}
```

You do not need to configure an IP address or a DLCI on this link.

2. To configure the logical address for the MLPPP, MLFR FRF.15, or MLFR FRF.16 bundle, include the **address** and **destination** statements:

```
address address {
  destination address;
}
```

You can include these statements at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number family inet]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet]`

When you add statements such as **mrru** to the configuration and commit, the T1 interface becomes part of the multilink bundle.



NOTE: For MLPPP and MLFR (FRF.15 and FRF.16) links, you must specify the subnet address as /32 or /30. Any other subnet designation is treated as a mismatch.

Related Documentation

- [Link and Multilink Services Overview on page 711](#)
- [Configuring the Number of Bundles on Link Services PICs on page 720](#)
- [Example: Configuring an MLPPP Bundle on page 750](#)

Example: Configuring a Link Services Interface with Two Links

This example uses the MLFR UNI NNI protocol between Router A and Router B and logically connects link services bundles **ls-1/1/0.3** and **ls-0/0/0.10**, as specified in [Table 29 on page 722](#).

Table 29: Link Services Bundle

Router A	Router B
t1-0/1/0 (ls-1/1/0:3)	t1-0/3/0 (ls-0/0/0:10)
t1-0/1/1 (ls-1/1/0:3)	t1-0/3/1 (ls-0/0/0:10)

For LMI to work properly, you must configure one router to be a DCE.

```
Configuration on [edit interfaces]
Router A         ls-1/1/0:3 {
                  dce;
                  encapsulation multilink-frame-relay-uni-nni;
                  unit 0 {
                    dlci 16;
                    family inet {
                      address 10.3.3.1/32 {
                        destination 10.3.3.2;
                      }
                    }
                  }
                }
              }
            t1-0/1/0 {
              encapsulation multilink-frame-relay-uni-nni;
              unit 0 {
                family mlfr-uni-nni {
                  bundle ls-1/1/0:3;
                }
              }
            }
            t1-0/1/1 {
              encapsulation multilink-frame-relay-uni-nni;
              unit 0 {
                family mlfr-uni-nni {
                  bundle ls-1/1/0:3;
                }
              }
            }
          }
```

```
Configuration on [edit interfaces]
Router B         ls-0/0/0:10 {
                  encapsulation multilink-frame-relay-uni-nni;
                  unit 0 {
                    dlci 16;
                    family inet {
                      address 10.3.3.2/32 {
                        destination 10.3.3.1;
                      }
                    }
                  }
                }
              }
            t1-0/3/0 {
              encapsulation multilink-frame-relay-uni-nni;
              unit 0 {
                family mlfr-uni-nni {
                  bundle ls-0/0/0:10;
                }
              }
            }
            t1-0/3/1 {
              encapsulation multilink-frame-relay-uni-nni;
              unit 0 {
                family mlfr-uni-nni {
```

```
        bundle ls-0/0/0:10;  
    }  
}  
}
```

- Related Documentation**
- [encapsulation \(Physical Interface\) on page 1604](#)
 - [Configuring Link Services Physical Interfaces on page 725](#)

Configuring the Physical and Logical Interfaces in a Multilink Configuration

- [Configuring Link Services Physical Interfaces on page 725](#)
- [Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 729](#)
- [Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces on page 730](#)
- [Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 732](#)
- [Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces on page 733](#)
- [Configuring MRRU on Multilink and Link Services Logical Interfaces on page 734](#)
- [Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces on page 735](#)
- [Example: Configuring Link Interfaces on Channelized MICs on page 735](#)

Configuring Link Services Physical Interfaces

You configure link services interface properties at the logical unit and physical interface level. Default settings for link services physical interface properties are described in “Default Settings for Link Services Interfaces” on page 726.

The following sections explain how to configure link services physical interfaces:

- [Default Settings for Link Services Interfaces on page 726](#)
- [Configuring Encapsulation for Link Services Physical Interfaces on page 726](#)
- [Configuring Acknowledgment Timers on Link Services Physical Interfaces on page 727](#)
- [Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16 on page 727](#)
- [Configuring Keepalives on Link Services Physical Interfaces on page 728](#)

For information about link services physical interface properties that can also be configured at the logical unit level, see “Multilink and Link Services Logical Interface Configuration Overview” on page 717.

Default Settings for Link Services Interfaces

Table 30 on page 726 lists the default settings for link services statements, together with the other permitted values or value ranges.

Table 30: Link Services Physical Interface Statements for MLFR FRF.16

Option	Default Value	Possible Values
Action red differential delay	remove-link	disable-tx, remove-link
Red differential delay	120 ms	1 through 2000 ms
Yellow differential delay	72 ms	1 through 2000 ms
Drop timeout period	0 ms	0 through 2000 ms
Encapsulation	multilink-frame-relay-uni-nni	multilink-frame-relay-uni-nni
Fragmentation threshold	0 bytes	128 through 16,320 bytes (Nx64)
LMI type	itu	ansi, itu
Minimum links	1 link	1 through 8 links
MRRU	1504 bytes	1500 through 4500 bytes
n391 (full status polling counter)	6	1 through 255
n392 (LMI error threshold)	3	1 through 10
n393 (LMI monitored event count)	4	1 through 10
t391 (link integrity verify polling timer)	10	5 through 30
t392 (polling verification timer)	15	5 through 30
Sequence ID format for MLFR	12 bits	12 bits

Configuring Encapsulation for Link Services Physical Interfaces

Link services interfaces support the physical interface encapsulation MLFR UNI NNI. By default, the physical interface encapsulation on link services interfaces is MLFR UNI NNI. Multilink interfaces do not support physical interface encapsulation.

For more information, see the *Junos OS Network Interfaces Library for Routing Devices*.

You can also configure logical interface encapsulation on multilink and link services interfaces. For more information, see “[Configuring Encapsulation for Multilink and Link Services Logical Interfaces](#)” on page 729.

To explicitly configure link services physical interface encapsulation, include the **encapsulation** statement at the **[edit interfaces ls-fpc/pic/port:channel]** hierarchy level:

encapsulation *type*;

You must also configure the T1, E1, or DS0 physical and physical interface with the same encapsulation type.

Configuring Acknowledgment Timers on Link Services Physical Interfaces

For link services interfaces configured with MLFR FRF.16, each link end point in a bundle initiates a request for bundle operation with its peer by transmitting an add link message. A hello message notifies the peer end point that the local end point is up. Both ends of a link generate a hello message periodically, or as configured with the hello timer. A remove link message notifies the peer that the local end management is removing the link from bundle operation. End points respond to add link, remove link, and hello messages by sending acknowledgment messages.

You can configure the maximum period to wait for an add link acknowledgment, hello acknowledgment, or remove link acknowledgment by including the **acknowledge-timer** statement at the **[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]** hierarchy level:

acknowledge-timer *milliseconds*;

The acknowledgment timer can be from 1 through 10 milliseconds. The default is 4 milliseconds.

For link services interfaces, you can configure the number of retransmission attempts to be made for consecutive hello or remove link messages after the expiration of the acknowledgment timer by including the **acknowledge-retries** statement at the **[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]** hierarchy level:

acknowledge-retries *number*;

acknowledgment-retries can be a value from 1 through 5. The default is 2.

You can configure the rate at which hello messages are sent by including the **hello-timer** statement at the **[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]** hierarchy level:

hello-timer *milliseconds*;

A hello message is transmitted after the specified period (in milliseconds) has elapsed. The hello timer can be from 1 through 180 milliseconds; the default is 10 milliseconds. When the hello timer expires, a link end point generates an add-link message.

Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16

For link services interfaces configured with MLFR FRF.16, the differential delay between links in a bundle is measured and warning is given when a link has a substantially greater differential delay than other links in the same bundle. The implementing endpoint can determine if the differential delay is in an acceptable range and decide to remove the link from the bundle, or to stop transmission on the link.

You can configure the yellow differential delay for links in a bundle by including the **yellow-differential-delay** statement at the **[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]** hierarchy level:

```
yellow-differential-delay milliseconds;
```

The yellow differential delay can be from 1 through 2000 milliseconds. The default is 72 milliseconds.

You can configure the red differential delay for links in a bundle to give warning by including the **red-differential-delay** statements at the **[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]** hierarchy level:

```
red-differential-delay milliseconds;
```

The red differential delay can be from 1 through 2000 milliseconds. The default is 120 milliseconds.

You can configure the action to be taken when differential delay exceeds the red limit by including the **action-red-differential-delay red** statements at the **[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]** hierarchy level:

```
action-red-differential-delay (disable-tx | remove-link);
```

The **disable-tx** option disables transmission on the link. The **remove-link** option removes the link from the bundle. The default action is **remove-link**.

You can view these settings in the output of the **show interfaces extensive lsq-fpc/pic/port:channel** command.

Configuring Keepalives on Link Services Physical Interfaces

You can tune the keepalive settings on the physical link-services interface. By default, the Junos OS uses ITU Q.933 Annex A LMIs for FRF.16. To instead use ITU Annex A LMIs (ANSI), include the **lmi-type ansi** statement at the **[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]** hierarchy level. LMI type ANSI is used in the following example:

```
lmi-type ansi;
```

To configure Frame Relay keepalive parameters on a link services interface, include the **n391**, **n392**, **n393**, **t391** and **t392** statements at the **[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]** hierarchy level:

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]  
n391 number;  
n392 number;  
n393 number;  
t391 number;  
t392 number;
```

The statements determine the indicated keepalive settings:

- **n391**—Full status polling interval. The data terminal equipment (DTE) sends a status inquiry to the data communication equipment (DCE) at the interval specified by the **t391** statement. This statement sets the frequency at which the DTE requests full status report; for example, the value **10** means that the DTE requests full status report in every tenth inquiry. The intermediate inquiries request a keepalive response only. The range is **1** through **255**, with a default of **6**.
- **n392**—Error threshold, which is the maximum number of errors that can occur during the number of events set by the **n393** statement before the link is marked inoperative. The range is **1** through **10**, with a default of **3**.
- **n393**—Monitored event count. The range is **1** through **10**, with a default of **4**.
- **t391**—The interval at which the DTE requests a keepalive response from the DCE and updates status, depending on the error threshold value. The range is **5** through **30** seconds, with a default of **10** seconds.
- **t392**—The period during which the DCE checks for keepalive responses from the DTE and updates status, depending on the DCE error threshold value. The range is from **5** through **30** seconds, with a default of **15** seconds.



NOTE: For the LMI to work properly, you must configure one side of a link services bundle to be a DCE.

Related Documentation

- [Link and Multilink Services Overview on page 711](#)
- [Example: Configuring a Link Services Interface with Two Links on page 722](#)

Configuring Encapsulation for Multilink and Link Services Logical Interfaces

Multilink and link services interfaces support the following logical interface encapsulation types:

- MLPPP
- MLFR end-to-end

By default, the logical interface encapsulation type on multilink interfaces is MLPPP. The default logical interface encapsulation type on link services interfaces is MLFR end-to-end. For general information on encapsulation, see the *Junos OS Network Interfaces Library for Routing Devices*.

You can also configure physical interface encapsulation on link services interfaces. For more information, see [“Configuring Encapsulation for Link Services Physical Interfaces” on page 726](#).

To configure multilink or link services encapsulation, include the **encapsulation** statement:

encapsulation *type*;

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

You must also configure the T1, E1, or DS0 physical interface with the same encapsulation type.



CAUTION: When you configure the first MLFR encapsulated unit or delete the last MLFR encapsulated unit on a port, it triggers an interface encapsulation change on the port, which causes an interface flap on the other units within the port that are configured with generic Frame Relay.

**Related
Documentation**

- [Link and Multilink Services Overview on page 711](#)
- [Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces on page 730](#)
- [Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 732](#)
- [Example: Configuring a Link Services Interface with MLPPP on page 753](#)
- [encapsulation \(Logical Interface\) on page 1603](#)

Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces

By default, the drop timeout parameter is disabled. You can configure a drop timeout value to provide a recovery mechanism if individual links in the multilink or link services bundle drop one or more packets. Drop timeout is not a differential delay tolerance setting, and does not limit the overall latency. However, you need to make sure the value you set is larger than the expected differential delay across the links, so that the timeout period does not elapse under normal jitter conditions, but only when there is actual packet loss. You can configure differential delay tolerance for link services interfaces only. For more information, see “[Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16](#)” on page 727.

To configure the drop timeout value, include the **drop-timeout** statement:

drop-timeout *milliseconds*;

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For link services interfaces, you also can configure the drop timeout value at the physical interface level by including the **drop-timeout** statement at the **[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]** hierarchy level:

drop-timeout *milliseconds*;

By default, the drop timer has a value of 500 ms for bundles greater than or equal to the T1 bandwidth value, and 1500 ms for other bundles. Any CLI-configured value overrides these defaults. Values can range from 1 through 2000 milliseconds. Values less than 5 milliseconds are not recommended, and a configured value of 0 reverts to the default value of 2000 milliseconds.



NOTE: For multilink or link services interfaces, if a packet or fragment encounters an error condition and is destined for a disabled bundle or link, it does not contribute to the dropped packet and frame counts in the per-bundle statistics. The packet is counted under the global error statistics and is not included in the global output bytes and output packet counts. This unusual accounting happens only if the error conditions are generated inside the multilink interface, not if the packet encounters errors on the wire or elsewhere in the network.

If you configure the **drop-timeout** statement with a value of 0, it disables any resequencing by the PIC for the specified class of MLPPP traffic. Packets are forwarded with the assumption that they arrived in sequence, and forwarding of fragmented packets is disabled for all classes. Fragments dropped as a result of this setting will increment the counter at the class level.

Alternatively, you can configure the **drop-timeout** statement at the **[edit class-of-service fragmentation-maps map-name forwarding-class class]** hierarchy level. The behavior and the default and range values are identical, but the setting applies only to the specified forwarding class. Configuration at the bundle level overrides configuration at the class-of-service level.

By default, compression of the inner PPP header in the MLPPP payload is enabled. To disable compression, include the **disable-mlppp-inner-ppp-pfc** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level. For example:

```
interfaces lsq-1/2/0 {
  unit 0 {
    encapsulation multilink-ppp;
    disable-mlppp-inner-ppp-pfc;
    multilink-max-classes 4;
    family inet {
      address 10.50.1.2/30;
    }
  }
}
```

For more information about CoS configuration, see the *Class of Service Feature Guide for Routing Devices*. You can view the configured drop-timeout value and the status of inner

PPP header compression by issuing the **show interfaces *interface-name* extensive** command.

**Related
Documentation**

- [Link and Multilink Services Overview on page 711](#)
- [Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 729](#)
- [Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 732](#)
- [Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces on page 733](#)

Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces

For multilink and link services logical interfaces with MLPPP encapsulation only, you can configure a *fragmentation threshold* to limit the size of packet payloads transmitted across the individual links within the multilink circuit. The software splits any incoming packet that exceeds the fragmentation threshold into smaller units suitable for the circuit size; it reassembles the fragments at the other end, but does not affect the output traffic stream. The threshold value affects the payload only; it does not affect the MLPPP header. By default, the fragmentation threshold parameter is disabled.



NOTE: To ensure proper load balancing:

- For Link Services MLFR (FRF.15 and FRF.16) interfaces, do not include the **fragment-threshold** statement in the configuration.
- For MLPPP interfaces, do not include both the **fragment-threshold** statement and the **short-sequence** statement in the configuration.
- For MLFR (FRF.15 and FRF.16) and MLPPP interfaces, if the MTU of links in a bundle is less than the bundle MTU plus encapsulation overhead, then fragmentation is automatically enabled. You should avoid this situation for MLFR (FRF.15 and FRF.16) interfaces and for MLPPP interfaces on which short-sequencing is enabled.

To configure a fragmentation threshold value, include the **fragment-threshold** statement:

fragment-threshold *bytes*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

For link services interfaces, you also can configure a fragmentation threshold value at the physical interface level by including the **fragment-threshold** statement at the **[edit interfaces *ls-fpc/pic/port:channel* mlfr-uni-nni-bundle-options]** hierarchy level:

fragment-threshold *bytes*;

The maximum fragment size can be from 128 through 16,320 bytes. The Junos OS automatically subdivides packet payloads that exceed this value. Any value you set must be a multiple of 64 bytes ($N \times 64$). The default value, 0, results in no fragmentation.

Related Documentation

- [Link and Multilink Services Overview on page 711](#)
- [Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 729](#)
- [Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces on page 730](#)
- [Example: Configuring a Multilink Interface with MLPPP on page 747](#)
- [fragment-threshold on page 1606](#)

Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces

You can set the minimum number of links that must be up for the multilink bundle as a whole to be labeled up. By default, only one link must be up for the bundle to be labeled up. A member link is considered up when the PPP Link Control Protocol (LCP) phase transitions to open state.

The **minimum-links** value should be identical on both ends of the bundle.

To set the minimum number, include the **minimum-links** statement:

minimum-links *number*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

For link services interfaces, you also can configure the minimum number of links at the physical interface level by including the **minimum-links** statement at the **[edit interfaces *ls-fpc/pic/port:channel* mlfr-uni-nni-bundle-options]** hierarchy level:

minimum-links *number*;

The number can be from 1 through 8. The maximum number of links supported in a bundle is 8. When 8 is specified, all configured links of a bundle must be up.

Related Documentation

- [Link and Multilink Services Overview on page 711](#)
- [Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 729](#)
- [Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces on page 730](#)
- [Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 732](#)

- [Configuring MRRU on Multilink and Link Services Logical Interfaces on page 734](#)

Configuring MRRU on Multilink and Link Services Logical Interfaces

The *maximum received reconstructed unit (MRRU)* is similar to a maximum transmission unit (MTU), but applies only to multilink bundles; it is the maximum packet size that the multilink interface can process. By default, the MRRU is set to 1500 bytes; you can configure a different MRRU value if the peer equipment allows this. The MRRU accounts for the original payload, for example the Layer 3 protocol payload, but does not include the 2-byte PPP header or the additional MLPPP or MLFR header applied while the individual multilink packets are traversing separate links in the bundle.

To configure a different MRRU value, include the **mrru** statement:

mrru bytes;

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For link services interfaces, you also can configure a different MRRU at the physical interface level by including the **mrru** statement at the [edit interfaces *ls-fpc/pic/port:channel* mlfr-uni-nni-bundle-options] hierarchy level:

mrru bytes;

The MRRU size can range from 1500 through 4500 bytes.



NOTE: If you set the MRRU on a bundle to a value larger than the MTU of the individual links within it, you must enable a fragmentation threshold for that bundle. Set the threshold to a value no larger than the smallest MTU of any link included in the bundle.

Determine the appropriate MTU size for the bundle by ensuring that the MTU size does not exceed the sum of the encapsulation overhead and the MTU sizes for the links in the bundle.

You can configure separate **family mtu** values on the following protocol families under bundle interfaces: **inet**, **inet6**, **iso**, and **mpls**. If not configured, the default value of 1500 is used on all except for **mpls** configurations, in which the value 1488 is used.



NOTE: The effective family MTU might be different from the MTU value specified for MLPPP configurations, because it is adjusted downward by the remote MRRU's constraints. The remote MRRU configuration is not supported on M120 routers.

- Related Documentation**
- [Link and Multilink Services Overview on page 711](#)
 - [Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 729](#)
 - [Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces on page 730](#)
 - [Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 732](#)
 - [Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces on page 733](#)

Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces

For MLPPP, the sequence header format is set to 24 bits by default. You can configure an alternative value of 12 bits, but 24 bits is considered the more robust value for most networks.

To configure a different sequence header value, include the **short-sequence** statement:

short-sequence;

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

For MLFR FRF.15, the sequence header format is set to 24 bits by default. This is the only valid option.

- Related Documentation**
- [Link and Multilink Services Overview on page 711](#)
 - [Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces on page 730](#)
 - [Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 732](#)
 - [Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces on page 733](#)
 - [Configuring MRRU on Multilink and Link Services Logical Interfaces on page 734](#)
 - [Configuring DLCIs on Link Services Logical Interfaces on page 777](#)

Example: Configuring Link Interfaces on Channelized MICs

- [Requirements on page 736](#)
- [Overview on page 736](#)

- [Configuration on 4-port Channelized SONET/SDH OC3/STM1 \(Multi-Rate\) MIC with SFP on page 737](#)
- [Verification on page 744](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.1 or later for MX240, MX480, and MX960 routers
- One MX240, MX480, or MX960 router

Overview

This example provides information about configuring the link interfaces on the following channelized MICs:

- 4-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-4CHOC3-2CHOC12)
- 8-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-8CHOC3-4CHOC12)
- 8-port Channelized DS3/E3 MIC (MIC-3D-8CHDS3-E3-B)

You need to first partition each port on the MICs to configure the link interfaces T1, T3, and DS, and then you configure the link interfaces for bundles. An MLPPP bundle involves "bundling" multiple T1/T3/DS interfaces into a single, logical interface that uses only one IP address. For more information about MLPPP bundles, see ["Understanding MLPPP Bundles and Link Fragmentation and Interleaving \(LFI\) on Serial Links" on page 719](#). Similarly, you can partition the ports to configure the MICs to the E1/E3 interfaces by setting the framing mode to SDH.

For more information about multilink-based protocols on MX240, MX480, and MX960 routers with Multiservices DPC, see ["Multilink Interfaces on Channelized MICs Overview" on page 715](#).

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.



.....

NOTE: You can set the values for each parameter according to your requirement. The values given in this example are for illustration purposes only.

.....

Configuration on 4-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP

To partition each port on the MIC and configure the link interfaces T1, T3, and DS on it and to configure the link interfaces for bundles, perform the following tasks:

- [Partitioning Ports on the Channelized MICs and Configuring the Link Interfaces T1, T3, and DS on page 738](#)
- [Configuring MLPPP, MLFR FRF.15, and MLFR FRF.16 on Link Interfaces for Bundles on page 740](#)
- [Results on page 741](#)

CLI Quick Configuration

To quickly configure synchronization on the aforementioned routers, copy the following commands, paste them in a text file, remove any line breaks, and then copy and paste the commands into the CLI.

[edit]

```
set interfaces coc12-5/2/0 partition 1 interface-type coc1
set interfaces coc12-5/2/0 partition 1 oc-slice 1
set interfaces coc12-5/2/0 partition 2 oc-slice 2 interface-type coc1
set interfaces coc12-5/2/0 partition 3 oc-slice 3 interface-type coc1
set interfaces coc1-5/2/0:1 no-partition interface-type t3
set interfaces coc1-5/2/0:3 no-partition interface-type t3
set interfaces coc1-5/2/0:2 partition 1 interface-type ct1
set interfaces coc1-5/2/0:2 partition 2 interface-type t1
set interfaces coc1-5/2/0:2 partition 3 interface-type ct1
set interfaces coc1-5/2/0:2 partition 4 interface-type t1
set interfaces ct1-5/2/0:2:1 partition 1 timeslots 1 interface-type ds
set interfaces ct1-5/2/0:2:1 partition 2 timeslots 2 interface-type ds
set interfaces ct3-2/0/0 no-partition interface-type t3
set interfaces ct3-2/0/0 partition 1 interface-type t1
set interfaces ct3-2/0/0 partition 2 interface-type t1
set interfaces ct3-2/0/0 partition 3 interface-type t1
set interfaces ct3-2/0/0 partition 4 interface-type ct1
set interfaces ct1-5/2/0:2:1 partition 1 timeslots 1 interface-type ds
set interfaces t1-5/2/0:2:2 unit 0 family mlppp bundle rlsq0.1
set interfaces ds-5/2/0:2:1:1 unit 0 family mlppp bundle rlsq0.0
set interfaces ds-5/2/0:2:1:2 unit 0 family mlppp bundle rlsq0.0
set interfaces t3-5/2/0:3 unit 0 family mlppp bundle rlsq0.2
set interfaces t3-5/2/0:1 unit 0 family mlppp bundle rlsq0.2
set interfaces t1-5/2/0:2:2 encapsulation multilink-frame-relay-uni-nni unit 0 family
mlfr-uni-nni bundle rlsq0:0
set interfaces t1-5/2/0:2:4 encapsulation multilink-frame-relay-uni-nni unit 0 family
mlfr-uni-nni bundle rlsq0:0
set interfaces ds-5/2/0:2:1:1 encapsulation multilink-frame-relay-uni-nni unit 0 family
mlfr-uni-nni bundle rlsq0:0
set interfaces ds-5/2/0:2:1:2 encapsulation multilink-frame-relay-uni-nni unit 0 family
mlfr-uni-nni bundle rlsq0:0
```

```
set interfaces t1-5/2/0:2:2 encapsulation frame-relay unit 0 dlci 10 family mlfr-end-to-end
bundle rlsq0.0
set interfaces t1-5/2/0:2:4 encapsulation frame-relay unit 0 dlci 11 family mlfr-end-to-end
bundle rlsq0.0
set interfaces ds-5/2/0:2:1:1 encapsulation frame-relay unit 0 dlci 10 family
mlfr-end-to-end bundle rlsq0.0
set interfaces ds-5/2/0:2:1:2 encapsulation frame-relay unit 0 dlci 11 family
mlfr-end-to-end bundle rlsq0.0
set interfaces t3-5/2/0:1 encapsulation frame-relay unit 0 dlci 11 family mlfr-end-to-end
bundle rlsq0.1
set interfaces t3-5/2/0:3 encapsulation frame-relay unit 0 dlci 10 family mlfr-end-to-end
bundle rlsq0.1
```

Partitioning Ports on the Channelized MICs and Configuring the Link Interfaces T1, T3, and DS

Step-by-Step Procedure

To partition each port on the channelized MICs:

1. Configure the **coc12-5/2/0** interface by setting the **partition** option to 1 with the sublevel interface type set to **coc1**.

```
[edit interfaces]
user@host# set coc12-5/2/0 partition 1 interface-type coc1
```
2. Configure the **coc12-5/2/0** interface with the OC-slice range (OC-slice range specifies the bandwidth size required for the interface type you are configuring) set to 1.

```
[edit interfaces]
user@host# set coc12-5/2/0 partition 1 oc-slice 1
```
3. Configure the **coc12-5/2/0** interface by setting the **partition** option to 2 with the sublevel interface type set to **coc1** and the OC-slice range set to 2.

```
[edit interfaces]
user@host# set coc12-5/2/0 partition 2 oc-slice 2 interface-type coc1
```
4. Configure the **coc12-5/2/0** interface by setting the **partition** option to 3 with the sublevel interface type set to **coc1** and the OC-slice range set to 3.

```
[edit interfaces]
user@host# set coc12-5/2/0 partition 3 oc-slice 3 interface-type coc1
```
5. Configure the **coc1-5/2/0:1** interface as a clear channel by setting the **no-partition** option for the sublevel interface type **t3**. (A clear channel consolidates the entire bandwidth of a channelized interface into a single unpartitioned stream that looks like a standard interface.)

```
[edit interfaces]
user@host# set coc1-5/2/0:1 no-partition interface-type t3
```
6. Configure the **coc1-5/2/0:3** interface as a clear channel by setting the **no-partition** option for the sublevel interface type **t3**.

```
[edit interfaces]
user@host# set coc1-5/2/0:3 no-partition interface-type t3
```

7. Configure the **coc1-5/2/0:2** interface by setting the **partition** option to 1 and 3 with the sublevel interface type set to **ct1**. Configure the **coc1-5/2/0:2** interface by setting the **partition** option to 2 and 4 with the sublevel interface type set to **t1**.

```
[edit interfaces]
user@host# set coc1-5/2/0:2 partition 1 interface-type ct1
user@host# set coc1-5/2/0:2 partition 2 interface-type t1
user@host# set coc1-5/2/0:2 partition 3 interface-type ct1
user@host# set coc1-5/2/0:2 partition 4 interface-type t1
```

8. Configure the **ct1-5/2/0:2:1** interface by setting the **partition** option to 1 and 2 with the sublevel interface type set to **ds**. Configure the time slots for the partitions.

```
[edit interfaces]
user@host# set ct1-5/2/0:2:1 partition 1 timeslots 1 interface-type ds
user@host# set ct1-5/2/0:2:1 partition 2 timeslots 2 interface-type ds
```

9. Configure a clear channel on the channelized interface **ct3-2/0/0** by setting the **no-partition** option to the sublevel interface type **t3** (a clear channel consolidates the entire bandwidth of a channelized interface into a single unpartitioned stream that looks like a standard interface).

```
[edit interfaces]
user@host# set ct3-2/0/0 no-partition interface-type t3
```

10. Configure a clear channel on the channelized interface **ct3-2/0/0** by setting the **partition** option to 1, 2, and 3 with the sublevel interface type **ds**. Configure the **ct3-2/0/0** interface by setting the **partition** option to 4 with the sublevel interface type **ct1**.

```
[edit interfaces]
user@host# set ct3-2/0/0 partition 1 interface-type t1
user@host# set ct3-2/0/0 partition 2 interface-type t1
user@host# set ct3-2/0/0 partition 3 interface-type t1
user@host# set ct3-2/0/0 partition 4 interface-type ct1
```

11. Configure the **ct1-2/0/0:4** interface by setting the **partition** option to 1 and 2 with the sublevel interface type set to **ds**. Configure the time slots for the partitions.

```
[edit interfaces]
user@host# set ct1-5/2/0:2:1 partition 1 timeslots 1 interface-type ds
user@host# set ct1-5/2/0:2:1 partition 2 timeslots 2 interface-type ds
```

Results Display the results of partitioning each port on the MIC and configuring the link interfaces T1, T3, and DS:

Results for CHOC12/3 interfaces

```
user@host# show interfaces
coc12-5/2/0 {
  partition 1 oc-slice 1 interface-type coc1;
  partition 2 oc-slice 2 interface-type coc1;
  partition 3 oc-slice 3 interface-type coc1;
}
coc1-5/2/0:1 {
  no-partition interface-type t3;
}
coc1-5/2/0:3 {
  no-partition interface-type t3;
```

```
}
coc1-5/2/0:2 {
  partition 1 interface-type ct1;
  partition 2 interface-type t1;
  partition 3 interface-type ct1;
  partition 4 interface-type t1;
}
```

**Results for CHDS3 MIC
interfaces**

```
user@host# show interfaces
ct1-5/2/0:2:1 {
  partition 1 timeslots 1 interface-type ds;
  partition 2 timeslots 2 interface-type ds;
}
ct3-2/0/0 {
  no-partition interface-type t3;
  partition 1 interface-type t1;
  partition 2 interface-type t1;
  partition 3 interface-type t1;
  partition 4 interface-type ct1;
}
ct1-2/0/0:4 {
  partition 1 timeslots 1 interface-type ds;
  partition 2 timeslots 2 interface-type ds;
}
```

Configuring MLPPP, MLFR FRF.15, and MLFR FRF.16 on Link Interfaces for Bundles

**Step-by-Step
Procedure**

To configure MLPPP, MLFR FRF.15, and MLFR FRF.16 on the link interfaces T1, T3, and DS for bundles:

1. Configure the MLPPP encapsulation on the T1 link interfaces **t1-5/2/0:2:2** and **t1-5/2/0:2:4**.

```
[edit interfaces]
user@host# set t1-5/2/0:2:2 unit 0 family mlppp bundle rlsq0.1
user@host# set t1-5/2/0:2:4 unit 0 family mlppp bundle rlsq0.1
```
2. Configure the MLPPP encapsulation on the DS link interfaces **ds-5/2/0:2:1:1** and **ds-5/2/0:2:1:2**.

```
[edit interfaces]
user@host# set ds-5/2/0:2:1:1 unit 0 family mlppp bundle rlsq0.0
user@host# set ds-5/2/0:2:1:2 unit 0 family mlppp bundle rlsq0.0
```
3. Configure the MLPPP encapsulation on the T3 link interfaces **t3-5/2/0:3** and **t3-5/2/0:1**.

```
[edit interfaces]
user@host# set t3-5/2/0:3 unit 0 family mlppp bundle rlsq0.2
user@host# set t3-5/2/0:1 unit 0 family mlppp bundle rlsq0.2
```
4. Configure the MLFR FRF.16 encapsulation on the T1 link interfaces **t1-5/2/0:2:2** and **t1-5/2/0:2:4**.

```
[edit interfaces]
user@host# set t1-5/2/0:2:2 encapsulation multilink-frame-relay-uni-nni unit 0
family mlfr-uni-nni bundle rlsq0:0
```

```
user@host# set t1-5/2/0:2:4 encapsulation multilink-frame-relay-uni-nni unit 0
family mlfr-uni-nni bundle rlsq0:0
```

5. Configure the MLFR FRF.16 encapsulation on the DS link interfaces **ds-5/2/0:2:1:1** and **ds-5/2/0:2:1:2**.

```
[edit interfaces]
user@host# set ds-5/2/0:2:1:1 encapsulation multilink-frame-relay-uni-nni unit 0
family mlfr-uni-nni bundle rlsq0:0
user@host# set ds-5/2/0:2:1:2 encapsulation multilink-frame-relay-uni-nni unit 0
family mlfr-uni-nni bundle rlsq0:0
```

6. Configure the MLFR FRF.15 encapsulation on the T1 link interfaces **t1-5/2/0:2:2** and **t1-5/2/0:2:4**.

```
[edit interfaces]
user@host# set t1-5/2/0:2:2 encapsulation frame-relay unit 0 dlci 10 family
mlfr-end-to-end bundle rlsq0:0
user@host# set t1-5/2/0:2:4 encapsulation frame-relay unit 0 dlci 11 family
mlfr-end-to-end bundle rlsq0:0
```

7. Configure the MLFR FRF.15 encapsulation on the DS link interfaces **ds-5/2/0:2:1:1** and **ds-5/2/0:2:1:2**.

```
[edit interfaces]
user@host# set ds-5/2/0:2:1:1 encapsulation frame-relay unit 0 dlci 10 family
mlfr-end-to-end bundle rlsq0:0
user@host# set ds-5/2/0:2:1:2 encapsulation frame-relay unit 0 dlci 11 family
mlfr-end-to-end bundle rlsq0:0
```

8. Configure the MLFR FRF.15 encapsulation on the T3 link interfaces **t3-5/2/0:1** and **t3-5/2/0:3**.

```
[edit interfaces]
user@host# set t3-5/2/0:1 encapsulation frame-relay unit 0 dlci 11 family
mlfr-end-to-end bundle rlsq0:1
user@host# set t3-5/2/0:3 encapsulation frame-relay unit 0 dlci 10 family
mlfr-end-to-end bundle rlsq0:1
```

Results

Display the results of the configuration of link interfaces for bundles:

MLPPP on T1 links

```
user@host# show interfaces
t1-5/2/0:2:2 {
  unit 0 {
    family mlppp {
      bundle rlsq0:1;
    }
  }
}
t1-5/2/0:2:4 {
  unit 0 {
    family mlppp {
      bundle rlsq0:1;
    }
  }
}
```

MLPPP on DS links

```
user@host# show interfaces
ds-5/2/0:2:1:1 {
  unit 0 {
    family mlppp {
      bundle rlsq0.0;
    }
  }
}
ds-5/2/0:2:1:2 {
  unit 0 {
    family mlppp {
      bundle rlsq0.0;
    }
  }
}
```

MLPPP on T3 links

```
user@host# show interfaces
t3-5/2/0:3 {
  unit 0 {
    family mlppp {
      bundle rlsq0.2;
    }
  }
}
t3-5/2/0:1 {
  unit 0 {
    family mlppp {
      bundle rlsq0.2;
    }
  }
}
```

MLFR FRF.15 on T1 links

```
user@host# show interfaces
t1-5/2/0:2:2 {
  encapsulation frame-relay;
  unit 0 {
    dlci 10;
    family mlfr-end-to-end {
      bundle rlsq0.0;
    }
  }
}
t1-5/2/0:2:4 {
  encapsulation frame-relay;
  unit 0 {
    dlci 11;
    family mlfr-end-to-end {
      bundle rlsq0.0;
    }
  }
}
```

MLFR FRF.15 on DS links

```
user@host# show interfaces
ds-5/2/0:2:1:1 {
  encapsulation frame-relay;
```



```

    unit 0 {
      dlci 10;
      family mlfr-end-to-end {
        bundle rlsq0.0;
      }
    }
  }
ds-5/2/0:2:1:2 {
  encapsulation frame-relay;
  unit 0 {
    dlci 11;
    family mlfr-end-to-end {
      bundle rlsq0.0;
    }
  }
}

```

**MLFR FRF.15 on T3
links**

```

user@host# show interfaces
t3-5/2/0:1 {
  encapsulation frame-relay;
  unit 0 {
    dlci 11;
    family mlfr-end-to-end {
      bundle rlsq0.1;
    }
  }
}
t3-5/2/0:3 {
  encapsulation frame-relay;
  unit 0 {
    dlci 10;
    family mlfr-end-to-end {
      bundle rlsq0.1;
    }
  }
}

```

**MLFR FRF.16 on T1
links**

```

user@host# show interfaces
t1-5/2/0:2:2 {
  encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
    family mlfr-uni-nni {
      bundle rlsq0.0;
    }
  }
}
t1-5/2/0:2:4 {
  encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
    family mlfr-uni-nni {
      bundle rlsq0.0;
    }
  }
}

```

```
MLFR FRF.16 on DS    user@host# show interfaces
links                ds-5/2/0:2:1:1 {
                      encapsulation multilink-frame-relay-uni-nni;
                      unit 0 {
                        family mlfr-uni-nni {
                          bundle rlsq0:0;
                        }
                      }
                      ds-5/2/0:2:1:2 {
                      encapsulation multilink-frame-relay-uni-nni;
                      unit 0 {
                        family mlfr-uni-nni {
                          bundle rlsq0:0;
                        }
                      }
                      }
```

Verification

Confirm that the configuration is working properly.

- [Verifying the MLPPP Bundle on page 744](#)
- [Verifying the MLFR FRF.15 Configuration on page 744](#)
- [Verifying the MLFR FRF.16 Configuration on page 744](#)

Verifying the MLPPP Bundle

Purpose Verify that the constituent links are added to the bundle correctly.

Action From operational mode, enter the **show interfaces lsq-fpc/pic/port** command.

Meaning The output displays the constituent links that are added to the bundle. For more information about the **show interfaces lsq-fpc/pic/port** operational command, see the [CLI Explorer](#).

Verifying the MLFR FRF.15 Configuration

Purpose Verify the MLFR FRF.15 configuration.

Action From operational mode, enter the **show interfaces lsq-fpc/pic/port** command.

Meaning The output displays the standard status information about the specified link services IQ interface. For more information about the **show interfaces lsq-fpc/pic/port** operational command, see the [CLI Explorer](#).

Verifying the MLFR FRF.16 Configuration

Purpose Verify the MLFR FRF.16 configuration.

Action From operational mode, enter the **show interfaces lsq-fpc/pic/port** command.

Meaning The output displays the standard status information about the specified link services IQ interface. For more information about the **show interfaces lsq-fpc/pic/port** operational command, see the [CLI Explorer](#).

Related Documentation

- [Link and Multilink Services Overview on page 711](#)
- [Multilink Interfaces on Channelized MICs Overview on page 715](#)
- [Example: Configuring an MLPPP Bundle on page 750](#)
- [Example: Configuring Multilink Frame Relay FRF.15 on page 783](#)
- [Example: Configuring Multilink Frame Relay FRF.16 on page 779](#)

Bundling Multiple PPP Links on a Single Link Using MLPPP

- [Example: Configuring a Multilink Interface with MLPPP on page 747](#)
- [Example: Configuring a Multilink Interface with MLPPP over ATM 2 Interfaces on page 748](#)
- [Example: Configuring an MLPPP Bundle on page 750](#)
- [Example: Configuring a Link Services Interface with MLPPP on page 753](#)
- [Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End \(FRF.15\) for WAN Interfaces on page 754](#)
- [Enabling MLPPP Link Fragmentation and Interleaving on page 769](#)
- [Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces on page 773](#)

Example: Configuring a Multilink Interface with MLPPP

```
[edit interfaces]
ml-1/0/0 {
  unit 1 {
    fragment-threshold 128;
    family inet {
      address 192.168.5.1/32 {
        destination 192.168.200.200;
      }
    }
  }
  unit 10 {
    family inet {
      address 10.1.1.3/32 {
        destination 10.1.1.2;
      }
    }
  }
}
t1-5/1/0 {
  unit 0 {
    family mlppp {
      bundle ml-1/0/0.1;
    }
  }
}
```

```
}
t1-5/1/1 {
  unit 0 {
    family mlppp {
      bundle ml-1/0/0.1;
    }
  }
}
t1-5/1/2 {
  unit 0 {
    family mlppp {
      bundle ml-1/0/0.1;
    }
  }
}
```

**Related
Documentation**

- [Link and Multilink Services Overview on page 711](#)
- [Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 732](#)
- [fragment-threshold on page 1606](#)

Example: Configuring a Multilink Interface with MLPPP over ATM 2 Interfaces

```
[edit interfaces]
at-0/0/0 {
  atm-options {
    pic-type atm2;
    vpi 10;
  }
  unit 0 {
    encapsulation atm-mlppp-llc;
    ppp-options {
      chap {
        access-profile pe-B-ppp-clients;
        local-name "pe-A-at-0/0/0";
      }
    }
    keepalive interval 5 up-count 6 down-count 4;
    vci 10.120;
    family mlppp {
      bundle ls-0/3/0.0;
    }
  }
}
at-0/0/1 {
  atm-options {
    pic-type atm2;
    vpi 11;
  }
  unit 1 {
    encapsulation atm-mlppp-llc;
    ppp-options {
      chap {
```

```

        access-profile pe-B-ppp-clients;
        local-name "pe-A-at-0/0/0";
    }
}
keepalive interval 5 up-count 6 down-count 4;
vci 11.120;
family mlppp {
    bundle ls-0/3/0.0;
}
}
}
at-1/2/3 {
    atm-options {
        pic-type atm2;
        vpi 12;
    }
}
unit 2 {
    encapsulation atm-mlppp-llc;
    ppp-options {
        chap {
            access-profile pe-B-ppp-clients;
            local-name "pe-A-at-0/0/0";
        }
    }
    keepalive interval 5 up-count 6 down-count 4;
    vci 12.120;
    family mlppp {
        bundle ls-0/3/0.0;
    }
}
}
...
ls-0/3/0 {
    encapsulation multilink-ppp;
    interleave-fragments;
    keepalive;
    unit 0 {
        mrru 4500;
        short-sequence;
        fragment-threshold 16320;
        drop-timeout 2000;
        encapsulation multilink-ppp;
        interleave-fragments;
        minimum-links 8;
        family inet {
            address 10.10.0.1/32 {
                destination 10.10.0.2;
            }
        }
        family iso;
        family inet6 {
            address 2001:DB8:0:1/32 {
                destination 2001:DB8:0:2;
            }
        }
    }
}
}

```

```
...
}
```

Related Documentation

- [Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 729](#)
- [Example: Configuring a Multilink Interface with MLPPP over ATM 2 Interfaces on page 748](#)
- [encapsulation \(Logical Interface\) on page 1603](#)

Example: Configuring an MLPPP Bundle

This example shows how to configure an MLPPP bundle to increase traffic bandwidth.

- [Requirements on page 750](#)
- [Overview on page 750](#)
- [Configuration on page 750](#)
- [Verification on page 753](#)


Requirements

Before you begin, you should have two MX Series routers (MX240, MX480, or MX960 routers) configured with at least two serial interfaces that communicate over serial links.

Overview

In this example, you create the MLPPP bundle `lsq-0/0/0.0` at the logical unit level of the link services interface `lsq-0/0/0` on the MX Series routers R0 and R1. You then add the two serial interfaces `se-1/0/0` and `se-1/0/1` as constituent links to the multilink bundle. In [Figure 27 on page 750](#), your company's branch office is connected to its main branch using routers R0 and R1. You transmit data and voice traffic on two low-speed 1-Mbps serial links. To increase bandwidth, you configure MLPPP and join the two serial links `se-1/0/0` and `se-1/0/1` into the multilink bundle `lsq-0/0/0.0`. Then you configure LFI and CoS on R0 and R1 to enable them to transmit voice packets ahead of data packets.

Figure 27: Configuring MLPPP and LFI on Serial Links



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
For device R0
set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.10/24
set interfaces se-1/0/0 unit 0 family mlppp bundle lsq-0/0/0.0
```



```

set interfaces se-1/0/1 unit 0 family mlppp bundle lsq-0/0/0.0
set interfaces se-1/0/0 serial-options clocking-mode dce clock-rate 2.0mhz
set interfaces se-1/0/1 serial-options clocking-mode dce clock-rate 2.0mhz

```

For device R1

```

set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.9/24
set interfaces se-1/0/0 unit 0 family mlppp bundle lsq-0/0/0.0
set interfaces se-1/0/1 unit 0 family mlppp bundle lsq-0/0/0.0

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure MLPPP bundle:

1. Create an interface on both the routers.

```

[edit]
user@host# edit interfaces lsq-0/0/0 unit 0

```
2. Configure a family inet and define the IP address on device R0.

```

[edit interfaces lsq-0/0/0 unit 0]
user@host# set family inet address 10.0.0.10/24

```
3. Configure a family inet and define the IP address on device R1.

```

[edit interfaces lsq-0/0/0 unit 0]
user@host# set family inet address 10.0.0.9/24

```
4. Specify the names of the constituent links to be added to the multilink bundle on both the routers.

```

[edit interfaces]
user@host# edit se-1/0/0 unit 0
user@host# set family mlppp bundle lsq-0/0/0.0
[edit interfaces]
user@host# edit se-1/0/1 unit 0
user@host# set family mlppp bundle lsq-0/0/0.0

```
5. Set the serial options to the same values for both interfaces on R0.



NOTE: R0 is set as a DCE device. The serial options are not set for interfaces on R1. You can set the serial options according to your network setup.

```

[edit interfaces]
user@host# set se-1/0/0 serial-options clocking-mode dce clock-rate 2.0mhz
user@host# set se-1/0/1 serial-options clocking-mode dce clock-rate 2.0mhz

```

Results From configuration mode, confirm your configuration by entering the **show interfaces lsq-0/0/0**, **show interfaces se-1/0/0**, and **show interfaces se-1/0/1** commands for R0 and R1. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
For device R0
[edit]
  user@host# show interfaces lsq-0/0/0
family inet {
  address 10.0.0.10/24;
}
[edit]
user@host# show interfaces se-1/0/0
  clocking-mode dce;
  clock-rate 2.0mhz;
}
  unit 0 {
    family mlppp {
      bundle lsq-0/0/0.0;
    }
  }
[edit]
  user@host# show interfaces se-1/0/1
serial-options {
  clocking-mode dce;
  clock-rate 2.0mhz;
}
  unit 0 {
    family mlppp {
      bundle lsq-0/0/0.0;
    }
  }
}

For device R1
[edit]
user@host# show interfaces lsq-0/0/0
  family inet {
    address 10.0.0.9/24;
  }
}
[edit]
  user@host# show interfaces se-1/0/0
  unit 0 {
    family mlppp {
      bundle lsq-0/0/0.0;
    }
  }
}
[edit]
  user@host# show interfaces se-1/0/1
  unit 0 {
    family mlppp {
      bundle lsq-0/0/0.0;
    }
  }
}
```

If you are done configuring the router, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly:

- [Verifying the MLPPP Bundle on page 753](#)

Verifying the MLPPP Bundle

Purpose Verify that the constituent links are added to the bundle correctly.

Action From operational mode, enter the **show interfaces lsq-0/0/0 statistics** command.

- Related Documentation**
- [Link and Multilink Services Overview on page 711](#)
 - [Multilink Interfaces on Channelized MICs Overview on page 715](#)
 - [Understanding MLPPP Bundles and Link Fragmentation and Interleaving \(LFI\) on Serial Links on page 719](#)
 - [Example: Configuring Link Interfaces on Channelized MICs on page 735](#)
 - [Example: Configuring Multilink Frame Relay FRF.15 on page 783](#)
 - [Example: Configuring Multilink Frame Relay FRF.16 on page 779](#)

Example: Configuring a Link Services Interface with MLPPP

```
[edit interfaces]
t1-0/0/0 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle ls-0/3/0.0;
    }
  }
}
t1-0/0/1 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle ls-0/3/0.0;
    }
  }
}
ls-0/3/0 {
  unit 0 {
    encapsulation multilink-ppp;
    family inet {
      address 10.16.1.2/32 {
        destination 10.16.1.1;
      }
    }
    family iso;
    family inet6 {
```

```
        address 2001:DB8:1:2/126;
    }
}
```

**Related
Documentation**

- [Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 729](#)
- [encapsulation \(Logical Interface\) on page 1603](#)

Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces

Inline Multilink PPP (MLPPP), Multilink Frame Relay (FRF.16), and Multilink Frame Relay End-to-End (FRF.15) for time-division multiplexing (TDM) WAN interfaces provide bundling services through the Packet Forwarding Engine without requiring a PIC or Dense Port Concentrator (DPC).

This example shows how to configure a Multilink PPP (MLPPP) bundle and Multilink Frame Relay End-to-End (FRF.15) for additional bandwidth, load balancing, and redundancy by aggregating low-speed links such as T1 (WAN interfaces).

- [Requirements on page 754](#)
- [Overview on page 754](#)
- [Configuration on page 755](#)
- [Verification on page 760](#)

Requirements

This example uses the following hardware and software components:

- Two MX Series Routers
- Junos OS Release 14.1 or later release

Before you begin, configure two MX Series routers (the MX240, MX480, or MX960) with at least two WAN interfaces that communicate over T1 links.

Overview

Traditionally, bundling services are used to bundle multiple low-speed links to create a higher bandwidth pipe. This combined bandwidth is available to traffic from all links and supports link fragmentation and interleaving (LFI) on the bundle, reducing high-priority packet transmission delay.

This support includes multiple links on the same bundle as well as multiclass extension for MLPPP. Through this service you can enable bundling services without additional DPC slots to support Service DPC and free up the slots for other MICs.

Configuring inline MLPPP for WAN interfaces benefits the following services:


- CE-PE link for Layer 3 VPN and DIA service with public switched telephone networks (PSTN)-based access networks

- PE-P link when PSTN is used for MPLS networks

In this example, to increase bandwidth, you configure MLPPP and join the T1 links into the multilink bundle. You aggregate T1 links to create the MLFR FRF.15 bundle on two MX Series routers, R0 and R1, and set the interface to **lsq-**. You configure logical units on the **lsq-** interface and set the family type to **inet** and an IP address. Then you configure an IP address for the multilink bundle on the unit level of the interface. You define the multilink bundle as an MLFR FRF.15 bundle by specifying the MLFR end-to-end encapsulation type. You specify the names of the constituent links to be added to the multilink bundle and set the encapsulation type to **frame-relay**. You then define Router R0 as a DCE device and Router R1 as a DTE device. You set the DLCI value (range is from 16 through 1022). Finally, you set the multilink bundle to **lsq-**.

Topology

Figure 28: Configuring Inline MLPPP and Multilink Frame Relay End-End (FRF.15) for WAN Interfaces



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R0

```
set chassis fpc 1 pic 0 multi-link-layer-2-inline
set interfaces lsq-1/0/0 unit 0 encapsulation multilink-ppp
set interfaces lsq-1/0/0 unit 0 family inet address 5.1.1.1/24
set interfaces lsq-1/0/0 unit 1 encapsulation multilink-frame-relay-end-to-end
set interfaces lsq-1/0/0 unit 1 family inet address 6.1.1.1/24
set interfaces t1-1/0/0:1 unit 0 family mlppp bundle lsq-1/0/0.0
set interfaces t1-1/0/0:2 unit 0 family mlppp bundle lsq-1/0/0.0
set interfaces t1-1/0/0:3 dce
set interfaces t1-1/0/0:4 dce
set interfaces t1-1/0/0:3 encapsulation frame-relay
set interfaces t1-1/0/0:4 encapsulation frame-relay
set interfaces t1-1/0/0:3 unit 0 dlci 1 family mlfr-end-to-end bundle lsq-1/0/0.1
set interfaces t1-1/0/0:4 unit 0 dlci 2 family mlfr-end-to-end bundle lsq-1/0/0.1
```

Device R1

```
set chassis fpc 2 pic 0 multi-link-layer-2-inline
set interfaces lsq-2/0/0 unit 0 encapsulation multilink-ppp
set interfaces lsq-2/0/0 unit 0 family inet address 5.1.1.2/24
set interfaces lsq-2/0/0 unit 1 encapsulation multilink-frame-relay-end-to-end
set interfaces lsq-2/0/0 unit 1 family inet address 6.1.1.2/24
```

```
set interfaces t1-2/0/0:1 unit 0 family mlppp bundle lsq-2/0/0.0
set interfaces t1-2/0/0:2 unit 0 family mlppp bundle lsq-2/0/0.0
set interfaces t1-2/0/0:3 encapsulation frame-relay
set interfaces t1-2/0/0:4 encapsulation frame-relay
set interfaces t1-2/0/0:3 unit 0 dlci 1 family mlfr-end-to-end bundle lsq-2/0/0.1
set interfaces t1-2/0/0:4 unit 0 dlci 2 family mlfr-end-to-end bundle lsq-2/0/0.1
```

To Configure Router R0

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces:

1. Enable inline Layer 2 bundling services.

```
[edit]
user@R0# set chassis fpc 1 pic 0 multi-link-layer-2-inline
```

2. Create the interface, specify a logical unit on the multilink bundle, and set the family type.

```
[edit]
user@R0# set interfaces lsq-1/0/0 unit 0 family inet address 5.1.1.1/24
```

3. Specify the *encapsulation* type as MLPPP.

```
[edit]
user@R0# set interfaces lsq-1/0/0 unit 0 encapsulation multilink-ppp
```

4. Create the interface, specify another logical unit on the multilink bundle, and set the family type.

```
[edit]
user@R0# set interfaces lsq-1/0/0 unit 1 family inet address 6.1.1.1/24
```

5. Specify another unit and define the multilink bundle as an MLFR FRF.15 bundle.

```
[edit]
user@R0# set interfaces lsq-1/0/0 unit 1 encapsulation
multilink-frame-relay-end-to-end
```

6. Specify the names of the constituent links to be added to the multilink bundle.

```
[edit]
user@R0# set interfaces t1-1/0/0:1 unit 0 family mlppp bundle lsq-1/0/0.0
user@R0# set interfaces t1-1/0/0:2 unit 0 family mlppp bundle lsq-1/0/0.0
```

7. Define the router as a DCE device.

```
[edit]
user@R0# set interfaces t1-1/0/0:3 dce
user@R0# set interfaces t1-1/0/0:4 dce
```

8. Specify the DLCI as well as the multilink bundle to which the interface is to be added.

```
[edit ]
user@R0# set interfaces t1-1/0/0:3 unit 0 dlci 1 family mlfr-end-to-end bundle
lsq-1/0/0.1
```

```
user@R0# set interfaces t1-1/0/0:4 unit 0 dlci 2 family mlfr-end-to-end bundle
lsq-1/0/0.1
```

9. Specify the names of the constituent links to be added to the multilink bundle.

```
[edit]
user@R0# set interfaces t1-1/0/0:3 encapsulation frame-relay
user@R0# set interfaces t1-1/0/0:4 encapsulation frame-relay
```

To Configure Router R1

Step-by-Step Procedure

To configure inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces:

1. Enable inline Layer 2 bundling services.

```
[edit]
user@R1# set chassis fpc 2pic 0 multi-link-layer-2-inline
```

2. Create the interface, specify a logical unit on the multilink bundle and set the family type.

```
[edit]
user@R1# set interfaces lsq-2/0/0 unit 0 family inet address 5.1.1.2/24
```

3. Specify the encapsulation type as MLPPP.

```
[edit]
user@R1# set interfaces lsq-2/0/0 unit 0 encapsulation multilink-ppp
```

4. Create the interface, specify another logical unit on the multilink bundle and set the family type.

```
[edit]
user@R1# set interfaces lsq-2/0/0 unit 1 family inet address 6.1.1.2/24
```

5. Specify another unit and define the multilink bundle as an MLFR FRF.15 bundle.

```
[edit]
user@R1# set interfaces lsq-2/0/0 unit 1 encapsulation
multilink-frame-relay-end-to-end
```

6. Specify the names of the constituent links to be added to the multilink bundle.

```
[edit]
user@R1# set interfaces t1-2/0/0:1 unit 0 family mlppp bundle lsq-2/0/0.0
user@R1# set interfaces t1-2/0/0:2 unit 0 family mlppp bundle lsq-2/0/0.0
```

7. Specify the DLCI as well as the multilink bundle to which the interface is to be added.

```
[edit ]
user@R1# set interfaces t1-2/0/0:3 unit 0 dlci 1 family mlfr-end-to-end bundle
lsq-2/0/0.1
user@R1# set interfaces t1-2/0/0:4 unit 0 dlci 2 family mlfr-end-to-end bundle
lsq-2/0/0.1
```

8. Specify the names of the constituent links to be added to the multilink bundle.

```
[edit]
user@R1# set interfaces t1-2/0/0:3 encapsulation frame-relay
user@R1# set interfaces t1-2/0/0:4 encapsulation frame-relay
```

Results

For Router R0, from configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces lsq-1/0/0**, **show interfaces t1-1/0/0:1**, **show interfaces t1-1/0/0:2**, **show interfaces t1-1/0/0:3**, and **show interfaces t1-1/0/0:4** commands.

For Router R1, from configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces lsq-2/0/0**, **show interfaces t1-2/0/0:1**, **show interfaces t1-2/0/0:2**, **show interfaces t1-2/0/0:3**, and **show interfaces t1-2/0/0:4** commands.

If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

For Router R0:

```
[edit]
user@R0# show chassis
fpc 1 {
  pic 0 {
    multi-link-layer-2-inline;
  }
}

[edit]
user@R0# show interfaces lsq-1/0/0
unit 0 {
  encapsulation multilink-ppp;
  family inet {
    address 5.1.1.1/24;
  }
}
unit 1 {
  encapsulation multilink-frame-relay-end-to-end;
  family inet {
    address 6.1.1.1/24;
  }
}

[edit]
user@R0# show interfaces t1-1/0/0:1
unit 0 {
  family mlppp {
    bundle lsq-1/0/0.0;
  }
}

[edit]
user@R0# show interfaces t1-1/0/0:2
unit 0 {
  family mlppp {
    bundle lsq-1/0/0.0;
  }
}

[edit]
user@R0# show interfaces t1-1/0/0:3
```



```

dce;
encapsulation frame-relay;
unit 0 {
    dlci 1;
    family mlfr-end-to-end {
        bundle lsq-1/0/0.1;
    }
}

[edit]
user@R0# show interfaces t1-1/0/0:4
dce;
encapsulation frame-relay;
unit 0 {
    dlci 2;
    family mlfr-end-to-end {
        bundle lsq-1/0/0.1;
    }
}

```

If you are done configuring the router, enter **commit** from configuration mode.

For Router R1:

```

[edit]
user@R1# show chassis
fpc 2 {
    pic 0 {
        multi-link-layer-2-inline;
    }
}

[edit]
user@R1# show interfaces lsq-2/0/0
unit 0 {
    encapsulation multilink-ppp;
    family inet {
        address 5.1.1.2/24;
    }
}
unit 1 {
    encapsulation multilink-frame-relay-end-to-end;
    family inet {
        address 6.1.1.2/24;
    }
}

[edit]
user@R1# show interfaces t1-2/0/0:1
unit 0 {
    family mlppp {
        bundle lsq-2/0/0.0;
    }
}

[edit]
user@R1# show interfaces t1-2/0/0:2
unit 0 {

```

```
family mlppp {
  bundle lsq-2/0/0.0;
}

[edit]
user@R1# show interfaces t1-2/0/0:3
encapsulation frame-relay;
unit 0 {
  dlci 1;
  family mlfr-end-to-end {
    bundle lsq-2/0/0.1;
  }
}

[edit]
user@R1# show interfaces t1-2/0/0:4
encapsulation frame-relay;
unit 0 {
  dlci 2;
  family mlfr-end-to-end {
    bundle lsq-2/0/0.1;
  }
}
```

If you are done configuring the router, enter **commit** from configuration mode.

Verification

Verifying the MLPPP Bundle and the MLFR FRF.15 Configuration

- Purpose** Verify that the constituent links are added to the bundle correctly.
- Action** From operational mode, run the **show interfaces lsq-1/0/0 extensive** command.

Sample Output

```
user@R0> show interfaces lsq-1/0/0:0 extensive
Physical interface: lsq-1/0/0, Enabled, Physical link is Up
Interface index: 292, SNMP ifIndex: 1065, Generation: 4986
Link-level type: LinkService, MTU: 1504
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Frame exceptions:
Oversized frames 0
Errored input frames 0
```

```

    Input on disabled link/bundle      0
    Output for disabled link/bundle    0
    Queuing drops                      0
    Buffering exceptions:
      Packet data buffer overflow      0
      Fragment data buffer overflow    0
    Assembly exceptions:
      Fragment timeout                 0
      Missing sequence number          0
      Out-of-order sequence number     0
      Out-of-range sequence number     0
    Hardware errors (sticky):
      Data memory error                0
      Control memory error             0
    Egress queues: 8 supported, 4 in use
    Queue counters:

```

	Queued packets	Transmitted packets	Dropped packets
0	0	0	0
1	0	0	0
2	0	0	0
3	0	0	0

```

    Queue number:
    0             Mapped forwarding classes
                  best-effort
    1             expedited-forwarding
    2             assured-forwarding
    3             network-control

    Logical interface lsq-1/0/0.0 (Index 327) (SNMP ifIndex 113518) (Generation
    6213)
    Flags: Hardware-Down Up Point-To-Point SNMP-Traps 0x4000 Encapsulation:
    Multilink-PPP
    Last flapped: 2014-04-24 04:37:39 PDT (00:08:50 ago)
    Bandwidth: 0
    Bundle links information:
      Active bundle links      0
      Removed bundle links     2
      Disabled bundle links    0
    Bundle options:
      MRRU                      1504
      Remote MRRU               N/A
      Drop timer period         32767
      Inner PPP Protocol field compression enabled
      Sequence number format    long (24 bits)
      Fragmentation threshold   0
      Links needed to sustain bundle 1
      Multilink classes         0
      Link layer overhead       4.0 %
    Multilink class 0 status:
      Received sequence number   0x0
      Transmit sequence number   0xffffffff
      Packet drops               0 (0 bytes)
      Fragment drops             0 (0 bytes)
      MRRU exceeded              0
      Fragment timeout           0
      Missing sequence number    0
      Out-of-order sequence number 0
      Out-of-range sequence number 0

```

```

    Packet data buffer overflow      0
    Fragment data buffer overflow    0
    Multilink class drop timeout    0 (ms)
Statistics      Frames      fps      Bytes      bps
Bundle:
Multilink:
  Input :      0      0      0      0
  Output:      0      0      0      0
Network:
  Input :      0      0      0      0
  Output:      0      0      0      0
IPV6 Transit Statistics      Packets      Bytes
Network:
  Input :      0      0
  Output:      0      0
Link:
t1-1/0/0:1.0
  Up time: 00:00:00
  Input :      0      0      0      0
  Output:      0      0      0      0
t1-1/0/0:2.0
  Up time: 00:00:00
  Input :      0      0      0      0
  Output:      0      0      0      0
Multilink detail statistics:
Bundle:
Fragments:
  Input :      0      0      0      0
  Output:      0      0      0      0
Non-fragments:
  Input :      0      0      0      0
  Output:      0      0      0      0
LFI:
  Input :      0      0      0      0
  Output:      0      0      0      0
NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
Protocol inet, MTU: 1500, Generation: 6263, Route table: 0
Flags: Sendbcst-pkt-to-re, Protocol-Down
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 5.1.1/24, Local: 5.1.1.1, Broadcast: Unspecified, Generation:
4211

Logical interface lsq-1/0/0.1 (Index 328) (SNMP ifIndex 113519) (Generation
6214)
Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-FR
Last flapped: 2014-04-24 04:46:00 PDT (00:00:29 ago)
Bandwidth: 3072kbps
Bundle links information:
  Active bundle links      2
  Removed bundle links     0
  Disabled bundle links    0
Bundle options:
  MRRU                      1504
  Drop timer period         32767
  Inner PPP Protocol field compression enabled
  Sequence number format    short (12 bits)
  Fragmentation threshold   0
  Links needed to sustain bundle 1
  Multilink classes         0
  Link layer overhead       4.0 %

```

```

Multilink class 0 status:
  Received sequence number      0x0
  Transmit sequence number      0xffffffff
  Packet drops                  0 (0 bytes)
  Fragment drops                0 (0 bytes)
  MRRU exceeded                 0
  Fragment timeout              0
  Missing sequence number       0
  Out-of-order sequence number  0
  Out-of-range sequence number  0
  Packet data buffer overflow    0
  Fragment data buffer overflow  0
  Multilink class drop timeout  0 (ms)
Statistics      Frames      fps      Bytes      bps
Bundle:
Multilink:
  Input :         0         0         0         0
  Output:         0         0         0         0
Network:
  Input :         0         0         0         0
  Output:         0         0         0         0
Link:
  t1-1/0/0:3.0
    Up time: 00:00:29
    Input :         0         0         0         0
    Output:         0         0         0         0
  t1-1/0/0:4.0
    Up time: 00:00:29
    Input :         0         0         0         0
    Output:         0         0         0         0
Multilink detail statistics:
Bundle:
Fragments:
  Input :         0         0         0         0
  Output:         0         0         0         0
Non-fragments:
  Input :         0         0         0         0
  Output:         0         0         0         0
LFI:
  Input :         0         0         0         0
  Output:         0         0         0         0
Protocol inet, MTU: 1500, Generation: 6264, Route table: 0
Flags: Sendbroadcast-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 6.1.1/24, Local: 6.1.1.1, Broadcast: Unspecified, Generation:
4213

```

From the operational mode, enter the **show interfaces lsq-2/0/0 extensive** command.

```

user@R1> show interfaces lsq-2/0/0 extensive
Physical interface: lsq-2/0/0, Enabled, Physical link is Up
Interface index: 262, SNMP ifIndex: 44421, Generation: 270
Encapsulation: Multilink-PPPLink-level type: LinkService, MTU: 1504
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
Last flapped   : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes :          0          0 bps
Output bytes :          0          0 bps
Input packets:          0          0 pps
Output packets:          0          0 pps

```

```

IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Frame exceptions:
  Oversized frames 0
  Errored input frames 0
  Input on disabled link/bundle 0
  Output for disabled link/bundle 0
  Queuing drops 0
Buffering exceptions:
  Packet data buffer overflow 0
  Fragment data buffer overflow 0
Assembly exceptions:
  Fragment timeout 0
  Missing sequence number 0
  Out-of-order sequence number 0
  Out-of-range sequence number 0
Hardware errors (sticky):
  Data memory error 0
  Control memory error 0
Egress queues: 8 supported, 4 in use
Queue counters:
  Queued packets  Transmitted packets  Dropped packets

  0                0                0                0
  1                0                0                0
  2                0                0                0
  3                0                0                0

Queue number:      Mapped forwarding classes
0                  best-effort
1                  expedited-forwarding
2                  assured-forwarding
3                  network-control

```

Logical interface lsq-2/0/0.0 (Index 354) (SNMP ifIndex 44422) (Generation 167)

```

Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation:
Multilink-PPPEncapsulation: Multilink-PPP
Last flapped: 2014-04-24 04:50:19 PDT (00:00:51 ago)
Bandwidth: 3072kbps
Bundle links information:
  Active bundle links 2
  Removed bundle links 0
  Disabled bundle links 0
Bundle options:
  MRRU 1504
  Remote MRRU 1504
  Drop timer period 32767
  Inner PPP Protocol field compression enabled
  Sequence number format long (24 bits)
  Fragmentation threshold 0
  Links needed to sustain bundle 1
  Multilink classes 0
  Link layer overhead 4.0 %
Multilink class 0 status:
  Received sequence number 0x0

```

```

    Transmit sequence number      0xffffffff
    Packet drops                  0 (0 bytes)
    Fragment drops                0 (0 bytes)
    MRRU exceeded                 0
    Fragment timeout              0
    Missing sequence number       0
    Out-of-order sequence number  0
    Out-of-range sequence number  0
    Packet data buffer overflow    0
    Fragment data buffer overflow  0
    Multilink class drop timeout  0 (ms)
Statistics      Frames      fps      Bytes      bps
Bundle:
Multilink:
  Input :          0          0          0          0
  Output:          0          0          0          0
Network:
  Input :          0          0          0          0
  Output:          0          0          0          0
IPV6 Transit Statistics      Packets      Bytes
Network:
  Input :              0              0
  Output:              0              0
Link:
  t1-2/0/0:1.0
    Up time: 00:00:51
    Input :          0          0          0          0
    Output:          0          0          0          0
  t1-2/0/0:2.0
    Up time: 00:00:48
    Input :          0          0          0          0
    Output:          0          0          0          0
Multilink detail statistics:
Bundle:
Fragments:
  Input :          0          0          0          0
  Output:          0          0          0          0
Non-fragments:
  Input :          0          0          0          0
  Output:          0          0          0          0
LFI:
  Input :          0          0          0          0
  Output:          0          0          0          0
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mp1s:
Not-configured
Protocol inet, MTU: 1500, Generation: 199, Route table: 0
Flags: Sendbroadcast-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 5.1.1/24, Local: 5.1.1.2, Broadcast: Unspecified, Generation:
153

Logical interface lsq-4/0/0.1 (Index 355) (SNMP ifIndex 44423) (Generation 168)

Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-FR
Last flapped: 2014-04-24 04:50:19 PDT (00:00:51 ago)
Bandwidth: 3072kbps
Bundle links information:
  Active bundle links      2
  Removed bundle links     0
  Disabled bundle links    0
Bundle options:

```

```

MRRU                                1504
Drop timer period                    32767
Inner PPP Protocol field compression enabled
Sequence number format              short (12 bits)
Fragmentation threshold              0
Links needed to sustain bundle       1
Multilink classes                    0
Link layer overhead                  4.0 %
Multilink class 0 status:
Received sequence number             0x0
Transmit sequence number             0xffffffff
Packet drops                         0 (0 bytes)
Fragment drops                       0 (0 bytes)
MRRU exceeded                        0
Fragment timeout                     0
Missing sequence number              0
Out-of-order sequence number         0
Out-of-range sequence number         0
Packet data buffer overflow          0
Fragment data buffer overflow         0
Multilink class drop timeout         0 (ms)
Statistics      Frames      fps      Bytes      bps
Bundle:
Multilink:
  Input :           0           0           0           0
  Output:           0           0           0           0
Network:
  Input :           0           0           0           0
  Output:           0           0           0           0
Link:
  t1-2/0/0:3.0
    Up time: 00:00:51
    Input :           0           0           0           0
    Output:           0           0           0           0
  t1-2/0/0:4.0
    Up time: 00:00:51
    Input :           0           0           0           0
    Output:           0           0           0           0
Multilink detail statistics:
Bundle:
Fragments:
  Input :           0           0           0           0
  Output:           0           0           0           0
Non-fragments:
  Input :           0           0           0           0
  Output:           0           0           0           0
LFI:
  Input :           0           0           0           0
  Output:           0           0           0           0
Protocol inet, MTU: 1500, Generation: 200, Route table: 0
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 6.1.1/24, Local: 6.1.1.2, Broadcast: Unspecified, Generation:
155

```

From operational mode, enter the **show interfaces lsq-1/0/0 statistics** command.

```

user@R0> show interfaces lsq-1/0/0 statistics
Physical interface: lsq-1/0/0, Enabled, Physical link is Up
Interface index: 292, SNMP ifIndex: 1065
Link-level type: LinkService, MTU: 1504
Device flags   : Present Running

```


Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
 Last flapped : Never
 Statistics last cleared: Never
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)

Logical interface lsq-1/0/0.0 (Index 327) (SNMP ifIndex 113518)
 Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
 Last flapped: 2014-04-24 04:50:19 PDT (00:01:59 ago)
 Bandwidth: 3072kbps

Bundle links information:

Active bundle links 2
 Removed bundle links 0
 Disabled bundle links 0

Statistics	Frames	fps	Bytes	bps
------------	--------	-----	-------	-----

Bundle:

Multilink:

Input :	0	0	0	0
Output:	0	0	0	0

Network:

Input :	0	0	0	0
Output:	0	0	0	0

IPv6 Transit Statistics	Packets	Bytes
-------------------------	---------	-------

Network:

Input :	0	0
Output:	0	0

Link:

t1-1/0/0:1.0

Up time: 00:01:59

Input :	0	0	0	0
Output:	0	0	0	0

t1-1/0/0:2.0

Up time: 00:01:56

Input :	0	0	0	0
Output:	0	0	0	0

NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls: Not-configured

Protocol inet, MTU: 1500

Flags: Sendbroadcast-pkt-to-re

Addresses, Flags: Is-Preferred Is-Primary

Destination: 5.1.1/24, Local: 5.1.1.1

Logical interface lsq-1/0/0.1 (Index 328) (SNMP ifIndex 113519)

Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-FR

Last flapped: 2014-04-24 04:50:29 PDT (00:01:49 ago)

Bandwidth: 3072kbps

Bundle links information:

Active bundle links 2
 Removed bundle links 0
 Disabled bundle links 0

Statistics	Frames	fps	Bytes	bps
------------	--------	-----	-------	-----

Bundle:

Multilink:

Input :	0	0	0	0
Output:	0	0	0	0

Network:

Input :	0	0	0	0
Output:	0	0	0	0

Link:

t1-1/0/0:3.0

Up time: 00:01:49

```

      Input :          0          0          0          0
      Output:          0          0          0          0
t1-1/0/0:4.0
  Up time: 00:01:49
      Input :          0          0          0          0
      Output:          0          0          0          0
Protocol inet, MTU: 1500
Flags: Sendbcast-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 6.1.1/24, Local: 6.1.1.1

```

From operational mode, enter the **show interfaces lsq-2/0/0 statistics** command.

```

user@R1> show interfaces lsq-2/0/0 statistics
Physical interface: lsq-2/0/0, Enabled, Physical link is Up
Interface index: 262, SNMP ifIndex: 44421
Link-level type: LinkService, MTU: 1504
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
Last flapped : Never
Statistics last cleared: Never
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)

Logical interface lsq-2/0/0.0 (Index 354) (SNMP ifIndex 44422)
Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
Last flapped: 2014-04-24 04:50:19 PDT (00:04:33 ago)
Bandwidth: 3072kbps
Bundle links information:
  Active bundle links      2
  Removed bundle links     0
  Disabled bundle links    0
Statistics               Frames      fps      Bytes      bps
Bundle:
Multilink:
  Input :                  0          0          0          0
  Output:                  0          0          0          0
Network:
  Input :                  0          0          0          0
  Output:                  0          0          0          0
IPV6 Transit Statistics      Packets      Bytes
Network:
  Input :                  0          0
  Output:                  0          0
Link:
t1-2/0/0:1.0
  Up time: 00:04:33
  Input :                  0          0          0          0
  Output:                  0          0          0          0
t1-2/0/0:2.0
  Up time: 00:04:30
  Input :                  0          0          0          0
  Output:                  0          0          0          0
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
Protocol inet, MTU: 1500
Flags: Sendbcast-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 5.1.1/24, Local: 5.1.1.2

Logical interface lsq-2/0/0.1 (Index 355) (SNMP ifIndex 44423)
Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-FR

```

```

Last flapped: 2014-04-24 04:50:19 PDT (00:04:33 ago)
Bandwidth: 3072kbps
Bundle links information:
  Active bundle links      2
  Removed bundle links    0
  Disabled bundle links    0
Statistics      Frames      fps      Bytes      bps
Bundle:
Multilink:
  Input :           0          0          0          0
  Output:           0          0          0          0
Network:
  Input :           0          0          0          0
  Output:           0          0          0          0
Link:
t1-2/0/0:3.0
  Up time: 00:04:33
  Input :           0          0          0          0
  Output:           0          0          0          0
t1-2/0/0:4.0
  Up time: 00:04:33
  Input :           0          0          0          0
  Output:           0          0          0          0
Protocol inet, MTU: 1500
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 6.1.1/24, Local: 6.1.1.2

```

- Related Documentation**
- [Inline MLPPP for WAN Interfaces Overview on page 559](#)
 - [Enabling MLPPP Link Fragmentation and Interleaving on page 769](#)
 - [Example: Configuring Multilink Frame Relay FRF.15 on page 783](#)
 - [Link and Multilink Services Interfaces Feature Guide for Routing Devices](#)
 - [mlfr-uni-nni-bundles-inline on page 1412](#)
 - [multi-link-layer-2-inline on page 1414](#)
 - [pic \(MX Series Routers\)](#)
 - [show interfaces \(Link Services IQ\) on page 1885](#)

Enabling MLPPP Link Fragmentation and Interleaving

MLPPP enables you to bundle multiple PPP links into a single multilink bundle. MLPPP bundle support on an inline LSQ interface is identical to a non-inline LSQ interface, because the configuration to enable fragmentation, link fragmentation and interleaving (LFI), and timeout is identical.

Priority scheduling on a multilink bundle determines the order in which an output interface transmits traffic from an output queue. The queues are serviced in a weighted round-robin fashion. But when a queue containing large packets starts using the multilink bundle, small and delay-sensitive packets must wait their turn for transmission. Because of this delay, some slow links, such as T1 and E1, can become useless for delay-sensitive traffic.

Link fragmentation and interleaving (LFI) solves this problem. It reduces delay and jitter on links by fragmenting large packets and interleaving delay-sensitive packets with the resulting smaller packets for simultaneous transmission across multiple links of a multilink bundle.

To configure schedule maps and fragmentation maps for MLPPP LFI:

1. Assign each forwarding class to an internal queue number by including the **forwarding-classes** statement at the **[edit class-of-service]** hierarchy level.

```
[edit class-of-service]
forwarding-classes {
  queue queue-number class-name;
}
```

For example, to set four output transmission queues:

```
[edit class-of-service]
user@host# set forwarding-classes queue 0 be
user@host# set forwarding-classes queue 1 ef
user@host# set forwarding-classes queue 2 af
user@host# set forwarding-classes queue 3 nc
```

2. To set a per-forwarding class fragmentation threshold, include the **fragment-threshold** statement in the **fragmentation-maps**.

```
[edit class-of-service]
fragmentation-maps {
  map-name{
    forwarding-class class-name {
      fragment-threshold bytes;
    }
  }
}
```

For example, to create two fragmentation maps and set a per-forwarding class fragmentation threshold:

```
[edit class-of-service]
user@host# set fragmentation-maps fragmap-1 forwarding-class af
  fragment-threshold 320
user@host# set fragmentation-maps fragmap-1 forwarding-class be
  fragment-threshold 256
user@host# set fragmentation-maps fragmap-1 forwarding-class ef
  fragment-threshold no-fragmentation
user@host# set fragmentation-maps fragmap-2 forwarding-class af
  fragment-threshold 192
user@host# set fragmentation-maps fragmap-2 forwarding-class be
  fragment-threshold 320
user@host# set fragmentation-maps fragmap-2 forwarding-class ef
  fragment-threshold 192
user@host# set fragmentation-maps fragmap-2 forwarding-class nc
  fragment-threshold no-fragmentation
```

The **fragment-threshold** statement in the LSQ bundle logical interface configuration applies to all forwarding classes. The **fragment-threshold** statement in **fragmentation-maps** for a particular **forwarding class**, if present, overrides the statement

configured in the LSQ bundle logical interface for that class. If **fragment-threshold** is not configured anywhere in the configuration, packets are still fragmented if **fragment-threshold** exceeds the smallest maximum transmission unit (MTU) or maximum received reconstructed unit (MRRU) of all links in the bundle.

3. Configure transmission scheduling parameters.

```
[edit class-of-service scheduler scheduler-name]
schedulers {
  scheduler-name {
    priority priority-level;
    transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
  }
}
```

For example, to set the transmit-rate percentage and the priority-level:

```
[edit class-of-service scheduler af-scheduler]
user@host# set transmit-rate percent 30
user@host# set priority low

[edit class-of-service scheduler be-scheduler]
user@host# set transmit-rate percent 20
user@host# set priority low

[edit class-of-service scheduler ef-scheduler]
user@host# set transmit-rate percent 35
user@host# set priority strict-high

[edit class-of-service scheduler nc-scheduler]
user@host# set transmit-rate percent 15
user@host# set priority high
```

4. After defining a scheduler, associate it with a specified forwarding class by including it in a *scheduler-map*.

```
[edit class-of-service]
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
```

For example, to associate the af-scheduler, be-scheduler, ef-scheduler, and nc-scheduler schedulers, with the af, be, ef, and nc forwarding-classes:

```
[edit class-of-service]
user@host# set scheduler-maps sched-map1 forwarding-class af scheduler
af-scheduler
user@host# set scheduler-maps sched-map1 forwarding-class be scheduler
be-scheduler
user@host# set scheduler-maps sched-map1 forwarding-class ef scheduler
ef-scheduler
user@host# set scheduler-maps sched-map1 forwarding-class nc scheduler
nc-scheduler
```

5. Configure traffic shaping and scheduling profiles.

```
[edit class-of-service]
traffic-control-profiles {
```

```
profile-name {  
    guaranteed-rate (percent percentage | rate) <burst-size bytes>;  
    scheduler-map map-name;  
    shaping-rate (percent percentage | rate) <burst-size bytes>;  
}  
}
```

For example, to set the traffic-control policies:

```
[edit class-of-service traffic-control-policies m1-tcp1]  
user@host# set guaranteed-rate 1m  
user@host# set scheduler-map sched-map1  
user@host# set shaping-rate 1m
```

6. Configure interface-specific CoS properties for incoming packets.

```
[edit class-of-service]  
interfaces {  
    interface-name {  
        unit logical-unit-number {  
            fragmentation-map map-name;  
            output-traffic-control-profile profile-name;  
        }  
    }  
}
```

For example, to apply the specified CoS traffic control profile (traffic scheduling and shaping configuration objects) to the output traffic at the logical interface:

```
[edit class-of-service]  
user@host# set interfaces lsq-0/1/0 unit 100 fragmentation-map fragmap-1  
output-traffic-control-profile m1-tcp1
```

The following partial configuration shows when the fragment threshold for low priority queues inherits from the fragment threshold configured in the bundle IFL and will have the value of 640.

```
[edit class-of-service]  
forwarding-classes {  
    queue 0 be;  
    queue 1 ef;  
    queue 2 af;  
    queue 3 nc;  
}  
fragmentation-maps {  
    fragmap-3 {  
        forwarding-class ef {  
            no-fragmentation;  
        }  
    }  
}  
schedulers {  
    af-scheduler {  
        transmit-rate percent 30;  
        priority low;  
    }  
    be-scheduler {  
        transmit-rate percent 20;
```

```
        priority low;
    }
    ef-scheduler {
        transmit-rate percent 35 rate-limit;
        priority strict-high;
    }
    nc-scheduler {
        transmit-rate percent 15;
        priority high;
    }
}
}
....
```

Related Documentation

- [Link and Multilink Services Interfaces Feature Guide for Routing Devices](#)
- [Understanding MLPPP Bundles and Link Fragmentation and Interleaving \(LFI\) on Serial Links on page 719](#)
- [Inline MLPPP for WAN Interfaces Overview on page 559](#)
- [Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End \(FRF.15\) for WAN Interfaces on page 754](#)
- [Example: Configuring Inline Multilink Frame Relay \(FRF.16\) for WAN Interfaces on page 788](#)

Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces

For link services FRF.15 and MLPPP interfaces only, you can configure link fragment interleaving (LFI). LFI reduces excessive delays of Frame Relay packets by fragmenting long packets into smaller packets and interleaving them with real-time frames. This allows real-time and non-real-time data frames to be carried together on lower-speed links without causing excessive delays to the real-time traffic. When the peer interface receives the smaller fragments, it reassembles the fragments into their original packet. For example, short delay-sensitive packets, such as packetized voice, can race ahead of larger delay-insensitive packets, such as common data packets.



NOTE: All Link Services PICs (4-multilink bundle, 32-multilink bundle, and 128-multilink bundle) support up to 256 link services interfaces with LFI enabled, if those link services interfaces contain only one constituent link each. For the Link Services PIC, multiple-link LFI bundles are simply multilink bundles, and are limited based on the type of PIC (4-multilink bundle, 32-multilink bundle, and 128-multilink bundle).

In addition, the multilink bundles you configure subtract from the total of 256 possible LFI-enabled link services interfaces. For example, if a 32-multilink bundle Link Services PIC has 24 multilink bundles configured and active, then you can configure $256 - 24 = 232$ LFI-enabled link services interfaces, each with a single constituent link.

For link services IQ interfaces (**lsq**), the **interleave-fragments** statement is not valid. Instead, you can enable LFI by configuring fragmentation maps. For more information, see [“Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces” on page 527](#).

You can configure multiple links in a bundle and configure packet interleaving. However, if you use packet interleaving, high-priority, nonmultilink-encapsulated packets use a hash-based algorithm to choose a single link.

For detailed information about link services CoS, see [“Configuring CoS on Link Services Interfaces” on page 801](#).

Per-bundle CoS queuing is supported on link services IQ interfaces (**lsq**). For more information about link services IQ interfaces, see [“Layer 2 Service Package Capabilities and Interfaces” on page 543](#).

The Junos OS supports end-to-end fragmentation in compliance with the FRF.12 *Frame Relay Fragmentation Implementation Agreement* standard. Unlike user-to-network interface (UNI) and network-to-network (NNI) fragmentation, end-to-end supports fragmentation only at the endpoints.

By default, packet interleaving is disabled. To enable packet interleaving, include the **interleave-fragments** statement:

interleave-fragments;

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Configuring LFI with DLCI Scheduling

For Link Services and Channelized DS3 IQ PICs, you can configure LFI and DLCI scheduling. For channelized DS3 interfaces, LFI is supported with FRF.15 only, and on M10i and M20 platforms only.

Configuring LFI with DLCI scheduling enables packets entering the Link Services PIC to be fragmented before being transmitted to the Channelized DS3 IQ PIC. Once the fragmented packets enter the Channelized DS3 IQ PIC, they are scheduled at the DLCI level, to allow priority transmission for real-time applications.

For more information about associating a scheduler with a DLCI, see the *Class of Service Feature Guide for Routing Devices*.

Example: Configuring LFI with DLCI Scheduling

Configure packets entering the Link Services PIC to be fragmented before being transmitted to the Channelized DS3 IQ PIC. Once the fragmented packets enter the Channelized DS3 IQ PIC, they are scheduled at the DLCI level, to allow priority transmission for real-time applications.

```
[edit interfaces]
ls-1/0/0 {
  unit 1 {
    encapsulation multilink-frame-relay-end-to-end;
    interleave-fragments;
    family inet {
      address 192.168.5.2/32 {
        destination 192.168.5.3;
      }
    }
  }
}
t3-1/0/0:1 {
  per-unit-scheduler;
  unit 0 {
    dlc1 16;
    encapsulation multilink-frame-relay-end-to-end;
    family mlfr-end-to-end {
      bundle ls-1/0/0.1;
    }
  }
}
[edit class-of-service]
interfaces {
  t3-1/0/0:1 {
    unit 0 {
      scheduler-map sched-map-logical-0;
      shaping-rate 10m;
    }
    unit 1 {
      scheduler-map sched-map-logical-1;
      shaping-rate 20m;
    }
  }
}
scheduler-maps {
  sched-map-logical-0 {
    forwarding-class best-effort scheduler sched-best-effort-0;
    forwarding-class assured-forwarding scheduler sched-bronze-0;
    forwarding-class expedited-forwarding scheduler sched-silver-0;
    forwarding-class network-control scheduler sched-gold-0;
```

```
}
sched-map-logical-1 {
  forwarding-class best-effort scheduler sched-best-effort-1;
  forwarding-class assured-forwarding scheduler sched-bronze-1;
  forwarding-class expedited-forwarding scheduler sched-silver-1;
  forwarding-class network-control scheduler sched-gold-1;
}
schedulers {
  sched-best-effort-0 {
    transmit-rate 4m;
  }
  sched-bronze-0 {
    transmit-rate 3m;
  }
  sched-silver-0 {
    transmit-rate 2m;
  }
  sched-gold-0 {
    transmit-rate 1m;
  }
  sched-best-effort-1 {
    transmit-rate 8m;
  }
  sched-bronze-1 {
    transmit-rate 6m;
  }
  sched-silver-1 {
    transmit-rate 4m;
  }
  sched-gold-1 {
    transmit-rate 2m;
  }
}
}
```

**Related
Documentation**

- [Link and Multilink Services Overview on page 711](#)
- [Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces on page 733](#)
- [Configuring MRRU on Multilink and Link Services Logical Interfaces on page 734](#)
- [Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces on page 735](#)
- [Configuring DLCIs on Link Services Logical Interfaces on page 777](#)

Bundling Multiple Frame Relay DLCIs into a Single Link Using MLFR

- [Configuring DLCIs on Link Services Logical Interfaces on page 777](#)
- [Example: Configuring a Multilink Interface with MLFR FRF.15 on page 778](#)
- [Example: Configuring Multilink Frame Relay FRF.16 on page 779](#)
- [Example: Configuring Multilink Frame Relay FRF.15 on page 783](#)
- [Example: Configuring a Link Services Interface with MLFR FRF.15 on page 786](#)
- [Example: Configuring a Link Services PIC with MLFR FRF.16 on page 787](#)
- [Example: Configuring Inline Multilink Frame Relay \(FRF.16\) for WAN Interfaces on page 788](#)

Configuring DLCIs on Link Services Logical Interfaces

For link services interfaces only, you can configure multiple DLCIs for MLFR FRF.16 or MLPPP bundles.

DLCIs are not supported on multilink interfaces.

Configuring Point-to-Point DLCIs for MLFR FRF.16 and MLPPP Bundles

For link services interfaces only, you can configure multiple point-to-point DLCIs for each MLFR FRF.16 or MLPPP bundle. A channelized interface, such as **ls-1/1/1:0**, denotes a single MLFR FRF.16 bundle. To configure a DLCI, include the **dlsi** statement:

dlsi *dlsi-identifier*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

The DLCI identifier is a value from 16 through 1022. Numbers 1 through 15 are reserved for future use.

When you configure point-to-point connections, the maximum transmission unit (MTU) sizes on both sides of the connection must be the same.

Configuring Multicast-Capable DLCIs for MLFR FRF.16 Bundles

For link services interfaces only, you can configure multiple multicast-capable DLCIs for each MLFR FRF.16 bundle. A channelized interface, such as **ls-1/1/1:0**, denotes a single MLFR FRF.16 bundle. By default, Frame Relay connections assume unicast traffic. If your Frame Relay switch performs multicast replication, you can configure the link services connection to support multicast traffic by including the **multicast-dlci** statement:

multicast-dlci *dlci-identifier*;

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The DLCI identifier is a value from 16 through 1022 that defines the Frame Relay DLCI over which the switch expects to receive multicast packets for replication.

You can configure multicast support only on point-to-multipoint link services connections. Multicast-capable DLCIs are not supported on multilink interfaces.

If keepalives are enabled, causing the interface to send Local Management Interface (LMI) messages during idle times, the number of possible DLCI configurations is limited by the MTU selected for the interface. For more information, see [“Configuring Keepalives on Link Services Physical Interfaces” on page 728](#).

Related Documentation

- [Link and Multilink Services Overview on page 711](#)
- [Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 732](#)
- [Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces on page 733](#)
- [Configuring MRRU on Multilink and Link Services Logical Interfaces on page 734](#)
- [Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces on page 735](#)

Example: Configuring a Multilink Interface with MLFR FRF.15

```
[edit interfaces]
ml-1/0/0 {
  unit 1 {
    encapsulation multilink-frame-relay-end-to-end;
    family inet {
      address 192.168.5.2/32 {
        destination 192.168.5.3;
      }
    }
  }
  unit 10 {
```

```

        encapsulation multilink-frame-relay-end-to-end;
        family inet {
            address 10.1.1.3/32 {
                destination 10.1.1.2;
            }
        }
    }
}
t1-5/1/0 {
    unit 0 {
        dlci 16;
        encapsulation multilink-frame-relay-end-to-end;
        family mlfr-end-to-end {
            bundle ml-1/0/0.1;
        }
    }
}
t1-5/1/1 {
    unit 0 {
        dlci 17;
        encapsulation multilink-frame-relay-end-to-end;
        family mlfr-end-to-end {
            bundle ml-1/0/0.10;
        }
    }
}
t1-5/1/2 {
    unit 0 {
        dlci 26;
        encapsulation multilink-frame-relay-end-to-end;
        family mlfr-end-to-end {
            bundle ml-1/0/0.10;
        }
    }
}
}

```

- Related Documentation**
- [Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 729](#)
 - [encapsulation \(Logical Interface\) on page 1603](#)

Example: Configuring Multilink Frame Relay FRF.16

This example shows how to configure MLFR FRF.16 for additional bandwidth, load balancing, and redundancy.

- [Requirements on page 780](#)
- [Overview on page 780](#)
- [Configuration on page 780](#)
- [Verification on page 783](#)

Requirements

Before you begin, you should have two MX Series 3D Universal Edge Routers configured with at least two serial interfaces that communicate over serial links.

Overview

In this example, you aggregate two T1 interfaces to create an MLFR FRF.16 bundle on two MX Series, R0 and R1. You configure the chassis interface and specify the number of MLFR FRF.16 bundles to be created on the interface. You then specify the channel to be configured as a multilink bundle and create interface lsq-0/0/0:0. You set the multilink bundle as an MLFR FRF.16 bundle by specifying the MLFR UNI NNI encapsulation type.

Then you define R0 as a DCE device and R1 as a DTE device. You configure a logical unit on the multilink bundle lsq-0/0/0:0, and set the family type to inet. You then assign a DLCI of 400 and an IP address of 10.0.0.10/24 to the multilink bundle. You create the T1 interfaces, t1-2/0/0 and t1-2/0/1, that are to be added as constituent links to the multilink bundle and define the Frame Relay encapsulation type. Finally, you set the multilink bundle to lsq-0/0/0:0.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
For device R0
set chassis fpc 0 pic 0 mlfr-uni-nni-bundles 1
set interfaces lsq-0/0/0:0 encapsulation multilink-frame-relay-uni-nni
set interfaces lsq-0/0/0 dce
set interfaces lsq-0/0/0 unit 0 dlci 400 family inet address 10.0.0.10/24
set interfaces t1-2/0/0 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/1 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/0 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
set interfaces t1-2/0/1 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
For device R1
set chassis fpc 0 pic 0 mlfr-uni-nni-bundles 1
set interfaces lsq-0/0/0:0 encapsulation multilink-frame-relay-uni-nni
set interfaces lsq-0/0/0 unit 0 dlci 400 family inet address 10.0.0.9/24
set interfaces t1-2/0/0 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/1 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/0 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
set interfaces t1-2/0/1 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an MLFR FRF.16 bundle:

1. Configure a chassis interface.

[edit]

- ```

user@host# edit chassis

```
2. Specify the number of MLFR bundles.
 

```

[edit chassis]
user@host# set fpc 0 pic 0 mlfr-uni-nni-bundles 1

```
  3. Create an interface.
 

```

[edit]
user@host# edit interfaces lsq-0/0/0:0

```
  4. Specify the MLFR encapsulation type.
 

```

[edit interfaces lsq-0/0/0:0]
user@host# set encapsulation multilink-frame-relay-uni-nni

```
  5. Set the router R0 as a DCE device.
 

```

[edit]
user@host# edit interfaces lsq-0/0/0
user@host# set dce

```
  6. Specify a logical unit on the multilink bundle and set the family type.
 

```

[edit interfaces lsq-0/0/0]
user@host# set unit 0 dlci 400 family inet address 10.0.0.10/24

```
  7. Create the T1 interfaces and set the Frame Relay encapsulation.
 

```

[edit interfaces]
user@host# set t1-2/0/0 encapsulation multilink-frame-relay-uni-nni
user@host# set t1-2/0/1 encapsulation multilink-frame-relay-uni-nni

```
  8. Specify the multilink bundle to which the interface is to be added as a constituent link on device R0.
 

```

[edit interfaces t1-2/0/0]
user@host# set unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0

```
  9. Specify the multilink bundle to which the interface is to be added as a constituent link on device R1.
 

```

[edit interfaces t1-2/0/1]
user@host# set unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0

```

**Results** From configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces lsq-0/0/0**, **show interfaces lsq-0/0/0:0**, **show interfaces t1-2/0/0**, and **show interfaces t1-2/0/1** commands for the routers R0 and R1. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

For device R0
[edit]
user@host# show chassis
fpc 0 {
pic 0 {
mlfr-uni-nni-bundles 1;
}
}
[edit]

```

```
user@host# show interfaces lsq-0/0/0
dce;
unit 0 {
dlci 400;
family inet {
address 10.0.0.10/24;
}
}
[edit]
user@host# show interfaces lsq-0/0/0:0
encapsulation multilink-frame-relay-uni-nni;
[edit]
user@host# show interfaces t1-2/0/0
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
family mlfr-uni-nni {
bundle lsq-0/0/0:0;
}
}
[edit]
user@host# show interfaces t1-2/0/1
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
family mlfr-uni-nni {
bundle lsq-0/0/0:0;
}
}

For device R1
[edit]
user@host# show chassis
unit 0 {
dlci 400;
family inet {
address 10.0.0.9/24;
}
}
[edit]
user@host# show interfaces lsq-0/0/0:0
encapsulation multilink-frame-relay-uni-nni;
[edit]
user@host# show interfaces t1-2/0/0
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
family mlfr-uni-nni {
bundle lsq-0/0/0:0;
}
}
[edit]
user@host# show interfaces t1-2/0/1
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
family mlfr-uni-nni {
bundle lsq-0/0/0:0;
}
}
}
```



If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly:

- [Verifying the MLFR FRF.16 Configuration on page 783](#)

---

### Verifying the MLFR FRF.16 Configuration

**Purpose** Verify the MLFR FRF.16 configuration.

**Action** From operational mode, enter the **show interfaces** command.

**Related Documentation**

- [Link and Multilink Services Overview on page 711](#)
- [Multilink Interfaces on Channelized MICs Overview on page 715](#)
- [Understanding MLPPP Bundles and Link Fragmentation and Interleaving \(LFI\) on Serial Links on page 719](#)
- [Example: Configuring Link Interfaces on Channelized MICs on page 735](#)
- [Example: Configuring an MLPPP Bundle on page 750](#)
- [Example: Configuring Multilink Frame Relay FRF.15 on page 783](#)

---

## Example: Configuring Multilink Frame Relay FRF.15

This example shows how to configure MLFR FRF.15 for additional bandwidth, load balancing, and redundancy by aggregating low-speed links such as T1, E1, and serial links.

- [Requirements on page 783](#)
- [Overview on page 783](#)
- [Configuration on page 784](#)
- [Verification on page 786](#)

## Requirements

Before you begin, you should have two MX Series 3D Universal Edge Routers (MX240, MX480, or MX960 routers) configured with at least two serial interfaces that communicate over serial links.

## Overview

In this example, you aggregate two T1 links to create the MLFR FRF.15 bundle on two MX Series routers, R0 and R1, and set the interface to `lsq-0/0/0`. You configure a logical unit on the `lsq-0/0/0` interface and set the family type to **inet** with address `10.0.0.4/24`. Then you configure an IP address for the multilink bundle on the unit level of the interface.

You define the multilink bundle as an MLFR FRF.15 bundle by specifying the MLFR end-to-end encapsulation type. You specify the names of the constituent links to be

added to the multilink bundle as t1-2/0/0 and t1-2/0/1 and set the encapsulation type to **frame-relay**. You then define R0 as a DCE device and R1 as a DTE device. You set the DLCI value to 100 (range is from 16 through 1022). Finally, you set the multilink bundle to lsq-0/0/0.0.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
For device R0
set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.4/24
set interfaces lsq-0/0/0 unit 0 encapsulation multilink-frame-relay-end-to-end
set interfaces t1-2/0/0 encapsulation frame-relay
set interfaces t1-2/0/1 encapsulation frame-relay
set interfaces lsq-0/0/0 dce
set interfaces lsq-0/0/0 unit 0 dlci 100 family mlfr-end-to-end bundle lsq-0/0/0.0
```

```
For device R1
set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.5/24
set interfaces lsq-0/0/0 unit 0 encapsulation multilink-frame-relay-end-to-end
set interfaces t1-2/0/0 encapsulation frame-relay
set interfaces t1-2/0/1 encapsulation frame-relay
set interfaces lsq-0/0/0 unit 0 dlci 100 family mlfr-end-to-end bundle lsq-0/0/0.0
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the MLFR FRF.15 bundle:

1. Create an interface on both the routers.  

```
[edit]
user@host# edit interfaces lsq-0/0/0 unit 0
```
2. Set a logical unit on the interface and define the family type for the routers R0 and R1.  

```
[edit interfaces lsq-0/0/0 unit 0]
user@host# set family inet address 10.0.0.4/24
user@host# set family inet address 10.0.0.5/24
```
3. Define the multilink bundle as an MLFR FRF.15 bundle.  

```
[edit interfaces lsq-0/0/0 unit 0]
user@host# set encapsulation multilink-frame-relay-end-to-end
```
4. Specify the names of the constituent links to be added to the multilink bundle.  

```
[edit interfaces]
user@host# set t1-2/0/0 encapsulation frame-relay
user@host# set t1-2/0/1 encapsulation frame-relay
```
5. Define the router R0 as a DCE device.  

```
[edit interfaces]
```

```
user@host# edit lsq-0/0/0
user@host# set dce
```

6. Specify the DLCI as well as the multilink bundle to which the interface is to be added.

```
[edit interfaces lsq-0/0/0]
user@host# set unit 0 dlci 100 family mlfr-end-to-end bundle lsq-0/0/0.0
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces lsq-0/0/0**, **show interfaces t1-2/0/0**, and **show interfaces t1-2/0/1** commands for R0 and R1. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For device R0

```
[edit]
user@host# show interfaces lsq-0/0/0
dce;
unit 0 {
 encapsulation multilink-frame-relay-end-to-end;
 dlci 100;
 family inet {
 address 10.0.0.4/24;
 }
 family mlfr-end-to-end {
 bundle lsq-0/0/0.0;
 }
}
[edit]
user@host# show interfaces t1-2/0/0
encapsulation frame-relay;
[edit]
user@host# show interfaces t1-2/0/1
encapsulation frame-relay;
```

For device R1

```
[edit]
user@host# show interfaces lsq-0/0/0
unit 0 {
 encapsulation multilink-frame-relay-end-to-end;
 dlci 100;
 family inet {
 address 10.0.0.5/24;
 }
 family mlfr-end-to-end {
 bundle lsq-0/0/0.0;
 }
}
[edit]
user@host# show interfaces t1-2/0/0
encapsulation frame-relay;
[edit]
user@host# show interfaces t1-2/0/1
encapsulation frame-relay;
```

If you are done configuring the router, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly:

- [Verifying the MLFR FRF.15 Configuration on page 786](#)

---

### Verifying the MLFR FRF.15 Configuration

---

**Purpose** Verify the MLFR FRF.15 configuration.

**Action** From operational mode, enter the **show interfaces** command.

**Related  
Documentation**

- [Link and Multilink Services Overview on page 711](#)
- [Multilink Interfaces on Channelized MICs Overview on page 715](#)
- [Understanding MLPPP Bundles and Link Fragmentation and Interleaving \(LFI\) on Serial Links on page 719](#)
- [Example: Configuring Link Interfaces on Channelized MICs on page 735](#)
- [Example: Configuring an MLPPP Bundle on page 750](#)
- [Example: Configuring Multilink Frame Relay FRF.16 on page 779](#)

---

## Example: Configuring a Link Services Interface with MLFR FRF.15

---

```
[edit interfaces]
t1-0/0/0 {
 encapsulation frame-relay;
 unit 0 {
 dlci 16;
 family mlfr-end-to-end {
 bundle ls-0/3/0.0;
 }
 }
}
t1-0/0/1 {
 encapsulation frame-relay;
 unit 0 {
 dlci 16;
 family mlfr-end-to-end {
 bundle ls-0/3/0.0;
 }
 }
}
ls-0/3/0 {
 unit 0 {
 encapsulation multilink-frame-relay-end-to-end;
 family inet {
 address 10.16.1.2/32 {
 destination 10.16.1.1;
 }
 }
 }
}
```

```

 family iso;
 family inet6 {
 address 2001:DB8:1:2/12;
 }
}
}

```

**Related Documentation**

- [encapsulation \(Logical Interface\) on page 1603](#)

## Example: Configuring a Link Services PIC with MLFR FRF.16

```

[edit chassis]
fpc 1 {
 pic 2 {
 mlfr-uni-nni-bundles 5;
 }
}
[edit interfaces]
t1-0/0/0 {
 encapsulation multilink-frame-relay-uni-nni;
 unit 0 {
 family mlfr-uni-nni {
 bundle ls-1/2/0:0;
 }
 }
}
t1-0/0/1 {
 encapsulation multilink-frame-relay-uni-nni;
 unit 0 {
 family mlfr-uni-nni {
 bundle ls-1/2/0:0;
 }
 }
}
ls-1/2/0:0 {
 dce;
 encapsulation multilink-frame-relay-uni-nni;
 unit 0 {
 dlci 26;
 family inet {
 address 10.26.1.1/32 {
 destination 10.26.1.2;
 }
 }
 }
}
}

```

**Related Documentation**

- [Configuring Link Services Physical Interfaces on page 725](#)
- [encapsulation \(Physical Interface\) on page 1604](#)

## Example: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces

---

Inline Multilink PPP (MLPPP), Multilink Frame Relay (FRF.16), and Multilink Frame Relay End-to-End (FRF.15) for time-division multiplexing (TDM) WAN interfaces provide bundling services through the Packet Forwarding Engine without requiring a PIC or Dense Port Concentrator (DPC).

Traditionally, bundling services are used to bundle multiple low-speed links to create a higher bandwidth pipe. This combined bandwidth is available to traffic from all links and supports link fragmentation and interleaving (LFI) on the bundle, reducing high priority packet transmission delay.

This support includes multiple links on the same bundle as well as multiclass extension for MLPPP. Through this service you can enable bundling services without additional DPC slots to support Service DPC and free up the slots for other MICs.

This example shows how to configure Multilink Frame Relay (FRF.16) for additional bandwidth, load balancing, and redundancy by aggregating low-speed links such as T1 (WAN interfaces).

- [Requirements on page 788](#)
- [Overview on page 788](#)
- [Configuration on page 789](#)
- [Verification on page 793](#)

### Requirements

This example uses the following hardware and software components:

- Two MX Series Routers
- Junos OS Release 14.1 or later release


Before you begin, configure two MX Series routers (the MX240, MX480, or MX960) with at least two WAN interfaces that communicate over T1 links.

### Overview

In this example, you aggregate T1 interfaces to create an MFR FRF.16 bundle on two MX Series routers, R0 and R1. You configure the chassis interface and specify the number of MFR FRF.16 bundles to be created on the interface. You then specify the channel to be configured as a multilink bundle and create interface **lsq-**. You set the multilink bundle as an MFR FRF.16 bundle by specifying the **multilink-frame-relay-uni-nni** encapsulation type. Then you define Router R0 as a DCE device and Router R1 as a DTE device. You configure a logical unit on the multilink bundle **lsq-**, and set the family type to **inet**. You create the T1 interfaces, that are to be added as constituent links to the multilink bundle and define the Frame Relay encapsulation type. Finally, you set the multilink bundle to **lsq-**.

## Topology

Figure 29: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces



## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R0

```
set chassis fpc 1 pic 0 mlfr-uni-nni-bundles-inline 1
set interfaces lsq-1/0/0:0 dce
set interfaces lsq-1/0/0:0 encapsulation multilink-frame-relay-uni-nni
set interfaces lsq-1/0/0:0 unit 0 dlci 10
set interfaces lsq-1/0/0:0 unit 1 dlci 20
set interfaces lsq-1/0/0:0 unit 0 family inet address 10.1.1.1/24
set interfaces lsq-1/0/0:0 unit 1 family inet address 11.1.1.1/24
set interfaces t1-1/0/0:5 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-1/0/0:6 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-1/0/0:5 unit 0 family mlfr-uni-nni bundle lsq-1/0/0:0
set interfaces t1-1/0/0:6 unit 0 family mlfr-uni-nni bundle lsq-1/0/0:0
```

Device R1

```
set chassis fpc 2 pic 0 mlfr-uni-nni-bundles-inline 1
set interfaces lsq-2/0/0:0 encapsulation multilink-frame-relay-uni-nni
set interfaces lsq-2/0/0:0 unit 0 dlci 10
set interfaces lsq-2/0/0:0 unit 1 dlci 20
set interfaces lsq-2/0/0:0 unit 0 family inet address 10.1.1.2/24
set interfaces lsq-2/0/0:0 unit 1 family inet address 11.1.1.2/24
set interfaces t1-2/0/0:5 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/0:6 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/0:5 unit 0 family mlfr-uni-nni bundle lsq-2/0/0:0
set interfaces t1-2/0/0:6 unit 0 family mlfr-uni-nni bundle lsq-2/0/0:0
```

### To Configure Router R0

---

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure inline Multilink Frame Relay (FRF.16) for WAN Interfaces:

1. Configure a chassis interface and specify the number of MFR bundles.  

```
[edit]
user@R0# set chassis fpc 1 pic 0 mlfr-uni-nni-bundles-inline 1
```
2. Create the interface and specify the MFR encapsulation type.  

```
[edit]
user@R0# set interfaces lsq-1/0/0:0 encapsulation multilink-frame-relay-uni-nni
```
3. Set Router R0 as a DCE device.  

```
[edit]
user@R0# set interfaces lsq-1/0/0:0 dce
```
4. Specify the DLCI value.  

```
[edit]
user@R0# set interfaces lsq-1/0/0:0 unit 0 dlci 10
user@R0# set interfaces lsq-1/0/0:0 unit 1 dlci 20
```
5. Specify a logical unit on the multilink bundle and set the family type.  

```
[edit]
user@R0# set interfaces lsq-1/0/0:0 unit 0 family inet address 10.1.1.1/24
user@R0# set interfaces lsq-1/0/0:0 unit 1 family inet address 11.1.1.1/24
```
6. Create the T1 interfaces and set the Frame Relay encapsulation.  

```
[edit]
user@R0# set interfaces t1-1/0/0:5 encapsulation multilink-frame-relay-uni-nni
user@R0# set interfaces t1-1/0/0:6 encapsulation multilink-frame-relay-uni-nni
```
7. Specify the multilink bundle to which the interface is to be added as a constituent link on Router R0.  

```
[edit]
user@R0# set interfaces t1-1/0/0:5 unit 0 family mlfr-uni-nni bundle lsq-1/0/0:0
user@R0# set interfaces t1-1/0/0:6 unit 0 family mlfr-uni-nni bundle lsq-1/0/0:0
```

### To Configure Router R1

---

**Step-by-Step Procedure** To configure inline Multilink Frame Relay (FRF.16) for WAN Interfaces:

1. Configure a chassis interface and specify the number of MFR bundles.  

```
[edit]
user@R1# set chassis fpc 2 pic 0 mlfr-uni-nni-bundles-inline 1
```
2. Create the interface and specify the MFR encapsulation type.  

```
[edit]
```



```
user@R1# set interfaces lsq-2/0/0:0 encapsulation multilink-frame-relay-uni-nni
```

3. Specify the DLCI value.

```
[edit]
user@R0# set interfaces lsq-2/0/0:0 unit 0 dlci 10
user@R0# set interfaces lsq-2/0/0:0 unit 1 dlci 20
```

4. Specify a logical unit on the multilink bundle and set the family type.

```
[edit]
user@R0# set interfaces lsq-2/0/0:0 unit 0 family inet address 10.1.1.2/24
user@R0# set interfaces lsq-2/0/0:0 unit 1 family inet address 11.1.1.2/24
```

5. Create the T1 interfaces and set the Frame Relay encapsulation.

```
[edit]
user@R1# set interfaces t1-2/0/0:5 encapsulation multilink-frame-relay-uni-nni
user@R1# set interfaces t1-2/0/0:6 encapsulation multilink-frame-relay-uni-nni
```

6. Specify the multilink bundle to which the interface is to be added as a constituent link on Router R1.

```
[edit]
user@R1# set interfaces t1-2/0/0:5 unit 0 family mlfr-uni-nni bundle lsq-2/0/0:0
user@R1# set interfaces t1-2/0/0:6 unit 0 family mlfr-uni-nni bundle lsq-2/0/0:0
```

## Results

For Router R0, from configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces lsq-1/0/0:0**, **show interfaces t1-1/0/0:5**, and **show interfaces t1-1/0/0:6** commands.

For Router R1, from configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces lsq-2/0/0:0**, **show interfaces t1-2/0/0:5**, and **show interfaces t1-2/0/0:6** commands.

If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

For Router R0:

```
[edit]
user@R0# show chassis
fpc 1 {
 pic 0 {
 mlfr-uni-nni-bundles-inline 1;
 }
}

[edit]
user@R0# show interfaces lsq-1/0/0:0
dce;
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
 dlci 10;
 family inet {
 address 10.1.1.1/24;
```

```
 }
 }
 unit 1 {
 dlci 20;
 family inet {
 address 11.1.1.1/24;
 }
 }
}

[edit]
user@R0# show interfaces t1-1/0/0:5
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
 family mlfr-uni-nni {
 bundle lsq-1/0/0:0;
 }
}

[edit]
user@R0# show interfaces t1-1/0/0:6
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
 family mlfr-uni-nni {
 bundle lsq-1/0/0:0;
 }
}
```

If you are done configuring the router, enter **commit** from configuration mode.

For Router R1:

```
[edit]
user@R1# show chassis
fpc 2 {
 pic 0 {
 mlfr-uni-nni-bundles-inline 1;
 }
}

[edit]
user@R1# show interfaces lsq-2/0/0:0
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
 dlci 10;
 family inet {
 address 10.1.1.2/24;
 }
}
unit 1 {
 dlci 20;
 family inet {
 address 11.1.1.2/24;
 }
}

[edit]
user@R1# show interfaces t1-2/0/0:5
encapsulation multilink-frame-relay-uni-nni;
```

```

unit 0 {
 family mlfr-uni-nni {
 bundle lsq-2/0/0:0;
 }
}

[edit]
user@R1# show interfaces t1-2/0/0:6
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
 family mlfr-uni-nni {
 bundle lsq-2/0/0:0;
 }
}

```

If you are done configuring the router, enter **commit** from configuration mode.

## Verification

### Verifying the MFR FRF.16 Configuration

**Purpose** Verify the MFR FRF.16 configuration.

**Action** From operational mode, run the **show interfaces lsq-1/0/0:0 extensive** command.

## Sample Output

```

user@R0> show interfaces lsq-1/0/0:0 extensive
Physical interface: lsq-1/0/0:0, Enabled, Physical link is Up
Interface index: 261, SNMP ifIndex: 122042, Generation: 4955
Link-level type: Multilink-FR-UNI-NNI, MTU: 1508
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
Last flapped : Never
Statistics last cleared: Never
Hold-times : Up 0 ms, Down 0 ms
Multilink Frame Relay UNI NNI bundle options:
 Device type DCE
 MRRU 1508
 Bandwidth 3072kbps
 Fragmentation threshold 0
 Red differential delay limit 120
 Yellow differential delay limit 72
 Red differential delay action Remove link
 Reassembly drop timer 65535
 Links needed to sustain bundle 1
 Link layer overhead 4.0 %
 LIP Hello timer 10
 Acknowledgement timer 4
 Acknowledgement retries 2
 Bundle class A
 LMI type Q.933 Annex A
 T391 LIV polling timer 10
 T392 polling verification timer 15
 N391 full status polling count 6
 N392 error threshold 3
 N393 monitored event count 4
Q.933 Annex A LMI settings: n392dce 3, n393dce 4, t392dce 15 seconds
LMI statistics:

```

```

Input : 52 (last seen 00:00:01 ago)
Output: 54 (last sent 00:00:01 ago)
DTE statistics:
 Enquiries sent : 0
 Full enquiries sent : 0
 Enquiry responses received : 0
 Full enquiry responses received : 0
DCE statistics:
 Enquiries received : 44
 Full enquiries received : 8
 Enquiry responses sent : 46
 Full enquiry responses sent : 8
Common statistics:
 Unknown messages received : 0
 Asynchronous updates received : 0
 Out-of-sequence packets received : 0
 Keepalive responses timedout : 1
Interface transmit statistics: Disabled
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Multilink Frame Relay UNI NNI bundle errors:
 Packet drops 0 (0 bytes)
 Fragment drops 0 (0 bytes)
 MRRU exceeded 0
 Exception events 0
Multilink Frame Relay UNI NNI bundle statistics:
 Frames fps Bytes bps

Multilink:
 Input : 0 0 0 0
 Output: 0 0 0 0
Network:
 Input : 0 0 0 0
 Output: 0 0 0 0
Multilink Frame Relay UNI NNI bundle links information:
 Active bundle links 2
 Removed bundle links 0
 Disabled bundle links 0
Multilink Frame Relay UNI NNI active bundle links statistics:
 Frames fps Bytes bps

t1-1/0/0:5
 Up time: 00:08:18
 Input : 0 0 0 0
 Output: 0 0 0 0
 Current differential delay 0.1 ms
 Recent high differential delay 0.8 ms
 Times over red diff delay 0
 Times over yellow diff delay 0
 LIP:add_lnk lnk_ack lnk_rej hello hel_ack lnk_rem rem_ack
 Rcv: 2 1 0 50 49 0 0
 Xmt: 16 2 0 49 50 1 0

t1-1/0/0:6
 Up time: 00:08:18

```

```

Input : 0 0 0 0
Output: 0 0 0 0
Current differential delay 0.0 ms
Recent high differential delay 0.7 ms
Times over red diff delay 0
Times over yellow diff delay 0
LIP:add_lnk lnk_ack lnk_rej hello hel_ack lnk_rem rem_ack
Rcv: 2 1 0 50 49 0 0
Xmt: 16 2 0 49 50 1 0

```

Logical interface lsq-1/0/0:0.0 (Index 336) (SNMP ifIndex 122044) (Generation 6209)

Flags: Up Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI

Last flapped: 2014-04-24 04:13:05 PDT (00:08:18 ago)

Multilink class 0 status:

```

Received sequence number 0x0
Transmit sequence number 0xffffffff

```

```

Packet drops 0 (0 bytes)
Fragment drops 0 (0 bytes)
MRRU exceeded 0
Fragment timeout 0
Missing sequence number 0
Out-of-order sequence number 0
Out-of-range sequence number 0
Packet data buffer overflow 0
Fragment data buffer overflow 0
Multilink class drop timeout 0 (ms)

```

| Statistics | Frames | fps | Bytes | bps |
|------------|--------|-----|-------|-----|
|------------|--------|-----|-------|-----|

Bundle:

Multilink:

```

Input : 0 0 0 0
Output: 0 0 0 0

```

Network:

```

Input : 0 0 0 0
Output: 0 0 0 0

```

Link:

t1-1/0/0:5

Up time: 00:08:18

```

Input : 0 0 0 0
Output: 0 0 0 0

```

t1-1/0/0:6

Up time: 00:08:18

```

Input : 0 0 0 0
Output: 0 0 0 0

```

Multilink detail statistics:

Bundle:

Fragments:

```

Input : 0 0 0 0
Output: 0 0 0 0

```

Non-fragments:

```

Input : 0 0 0 0
Output: 0 0 0 0

```

Protocol inet, MTU: 1500, Generation: 6258, Route table: 0

Flags: Sendbcst-pkt-to-re

Addresses, Flags: Is-Preferred Is-Primary

Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: Unspecified,

Generation: 4209

DLCI 10

Flags: Active

Total down time: 01:15:17 sec, Last down: 01:23:28 ago

Traffic statistics:

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

Logical interface lsq-1/0/0:0.1 (Index 337) (SNMP ifIndex 122067) (Generation 6210)

Flags: Up Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI

Last flapped: 2014-04-24 04:13:05 PDT (00:08:18 ago)

Multilink class 0 status:

```

Received sequence number 0x0
Transmit sequence number 0xffffffff
Packet drops 0 (0 bytes)
Fragment drops 0 (0 bytes)
MRRU exceeded 0
Fragment timeout 0
Missing sequence number 0
Out-of-order sequence number 0

```

Out-of-range sequence number 0

```

Packet data buffer overflow 0
Fragment data buffer overflow 0
Multilink class drop timeout 0 (ms)

```

| Statistics | Frames | fps | Bytes | bps |
|------------|--------|-----|-------|-----|
|------------|--------|-----|-------|-----|

Bundle:

Multilink:

|         |   |   |   |   |
|---------|---|---|---|---|
| Input : | 0 | 0 | 0 | 0 |
| Output: | 0 | 0 | 0 | 0 |

Network:

|         |   |   |   |   |
|---------|---|---|---|---|
| Input : | 0 | 0 | 0 | 0 |
| Output: | 0 | 0 | 0 | 0 |

Link:

t1-1/0/0:5

Up time: 00:08:18

|         |   |   |   |   |
|---------|---|---|---|---|
| Input : | 0 | 0 | 0 | 0 |
| Output: | 0 | 0 | 0 | 0 |

t1-1/0/0:6

Up time: 00:08:18

|         |   |   |   |   |
|---------|---|---|---|---|
| Input : | 0 | 0 | 0 | 0 |
| Output: | 0 | 0 | 0 | 0 |

Multilink detail statistics:

Bundle:

Fragments:

|         |   |   |   |   |
|---------|---|---|---|---|
| Input : | 0 | 0 | 0 | 0 |
| Output: | 0 | 0 | 0 | 0 |

Non-fragments:

|         |   |   |   |   |
|---------|---|---|---|---|
| Input : | 0 | 0 | 0 | 0 |
| Output: | 0 | 0 | 0 | 0 |

Protocol inet, MTU: 1500, Generation: 6260, Route table: 0

Flags: Sendbcst-pkt-to-re

Addresses, Flags: Is-Preferred Is-Primary

Destination: 11.1.1/24, Local: 11.1.1.1, Broadcast: Unspecified,

Generation: 4207

DLCI 20

Flags: Active

Total down time: 01:15:17 sec, Last down: 01:23:28 ago

Traffic statistics:

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

```
DLCI statistics:
 Active DLCI :2 Inactive DLCI :0
```

From the operational mode, enter the **show interfaces lsq-2/0/0:0 extensive** command.

```
user@R1> show interfaces lsq-2/0/0:0 extensive
Physical interface: lsq-2/0/0:0, Enabled, Physical link is Up
 Interface index: 232, SNMP ifIndex: 44389, Generation: 235
 Link-level type: Multilink-FR-UNI-NNI, MTU: 1508
 Device flags : Present Running
 Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
 Last flapped : Never
 Statistics last cleared: Never
 Hold-times : Up 0 ms, Down 0 ms
 Multilink Frame Relay UNI NNI bundle options:
 Device type DTE
 MRRU 1508
 Bandwidth 3072kbps
 Fragmentation threshold 0
 Red differential delay limit 120
 Yellow differential delay limit 72
 Red differential delay action Remove link
 Reassembly drop timer 65535
 Links needed to sustain bundle 1
 Link layer overhead 4.0 %
 LIP Hello timer 10
 Acknowledgement timer 4
 Acknowledgement retries 2
 Bundle class A
 LMI type Q.933 Annex A
 T391 LIV polling timer 10
 T392 polling verification timer 15
 N391 full status polling count 6
 N392 error threshold 3
 N393 monitored event count 4
 Q.933 Annex A LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds

 LMI statistics:
 Input : 80 (last seen 00:00:10 ago)
 Output: 100 (last sent 00:00:10 ago)
 DTE statistics:
 Enquiries sent : 82
 Full enquiries sent : 16
 Enquiry responses received : 67
 Full enquiry responses received : 13
 DCE statistics:
 Enquiries received : 0
 Full enquiries received : 0
 Enquiry responses sent : 0
 Full enquiry responses sent : 0
 Common statistics:
 Unknown messages received : 0
 Asynchronous updates received : 0
 Out-of-sequence packets received : 0
 Keepalive responses timeout : 1
 Interface transmit statistics: Disabled
 Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
```

## IPv6 transit statistics:

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

## Multilink Frame Relay UNI NNI bundle errors:

```

Packet drops 0 (0 bytes)
Fragment drops 0 (0 bytes)
MRRU exceeded 0
Exception events 0

```

## Multilink Frame Relay UNI NNI bundle statistics:

|            | Frames | fps | Bytes | bps |
|------------|--------|-----|-------|-----|
| Multilink: |        |     |       |     |
| Input :    | 0      | 0   | 0     | 0   |
| Output:    | 0      | 0   | 0     | 0   |
| Network:   |        |     |       |     |
| Input :    | 0      | 0   | 0     | 0   |
| Output:    | 0      | 0   | 0     | 0   |

## Multilink Frame Relay UNI NNI bundle links information:

```

Active bundle links 2
Removed bundle links 0
Disabled bundle links 0

```

## Multilink Frame Relay UNI NNI active bundle links statistics:

|                                | Frames | fps    | Bytes                         | bps |
|--------------------------------|--------|--------|-------------------------------|-----|
| t1-2/0/0:5                     |        |        |                               |     |
| Up time: 00:12:57              |        |        |                               |     |
| Input :                        | 0      | 0      | 0                             | 0   |
| Output:                        | 0      | 0      | 0                             | 0   |
| Current differential delay     |        | 0.0 ms |                               |     |
| Recent high differential delay |        | 2.8 ms |                               |     |
| Times over red diff delay      |        | 0      |                               |     |
| Times over yellow diff delay   |        | 0      |                               |     |
| LIP:add_lnk lnk_ack lnk_rej    |        |        | hello hel_ack lnk_rem rem_ack |     |
| Rcv: 1 2 0                     |        |        | 77 78 0 0                     |     |
| Xmt: 14 1 0                    |        |        | 78 77 0 0                     |     |
| t1-2/0/0:6                     |        |        |                               |     |
| Up time: 00:12:57              |        |        |                               |     |
| Input :                        | 0      | 0      | 0                             | 0   |
| Output:                        | 0      | 0      | 0                             | 0   |
| Current differential delay     |        | 0.0 ms |                               |     |
| Recent high differential delay |        | 2.8 ms |                               |     |
| Times over red diff delay      |        | 0      |                               |     |
| Times over yellow diff delay   |        | 0      |                               |     |
| LIP:add_lnk lnk_ack lnk_rej    |        |        | hello hel_ack lnk_rem rem_ack |     |
| Rcv: 1 2 0                     |        |        | 77 78 0 0                     |     |
| Xmt: 14 1 0                    |        |        | 78 77 0 0                     |     |

Logical interface lsq-2/0/0:0.0 (Index 348) (SNMP ifIndex 44399) (Generation 161)

Flags: Up Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI

Last flapped: 2014-04-24 04:13:05 PDT (00:12:57 ago)

## Multilink class 0 status:

```

Received sequence number 0x0
Transmit sequence number 0xffffffff
Packet drops 0 (0 bytes)
Fragment drops 0 (0 bytes)
MRRU exceeded 0
Fragment timeout 0
Missing sequence number 0
Out-of-order sequence number 0

```



```

Out-of-range sequence number 0
Packet data buffer overflow 0
Fragment data buffer overflow 0
Multilink class drop timeout 0 (ms)
Statistics Frames fps Bytes bps
Bundle:
Multilink:
 Input : 0 0 0 0
 Output: 0 0 0 0
Network:
 Input : 0 0 0 0
 Output: 0 0 0 0
Link:
t1-2/0/0:5
 Up time: 00:12:57
 Input : 0 0 0 0
 Output: 0 0 0 0
t1-2/0/0:6
 Up time: 00:12:57
 Input : 0 0 0 0
 Output: 0 0 0 0
Multilink detail statistics:
Bundle:
Fragments:
 Input : 0 0 0 0
 Output: 0 0 0 0
Non-fragments:
 Input : 0 0 0 0
 Output: 0 0 0 0
Protocol inet, MTU: 1500, Generation: 193, Route table: 0
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.1.1/24, Local: 10.1.1.2, Broadcast: Unspecified,
Generation: 149
DLCI 10
Flags: Active, DCE-Configured
Total down time: 00:03:18 sec, Last down: 00:15:38 ago
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Logical interface lsq-2/0/0:0.1 (Index 349) (SNMP ifIndex 44400) (Generation
162)
Flags: Up Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI
Last flapped: 2014-04-24 04:13:05 PDT (00:12:57 ago)
Multilink class 0 status:
Received sequence number 0x0
Transmit sequence number 0xffffffff
Packet drops 0 (0 bytes)
Fragment drops 0 (0 bytes)
MRRU exceeded 0
Fragment timeout 0
Missing sequence number 0
Out-of-order sequence number 0
Out-of-range sequence number 0
Packet data buffer overflow 0
Fragment data buffer overflow 0
Multilink class drop timeout 0 (ms)
Statistics Frames fps Bytes bps

```

```

Bundle:
 Multilink:
 Input : 0 0 0 0
 Output: 0 0 0 0
 Network:
 Input : 0 0 0 0
 Output: 0 0 0 0
Link:
 t1-2/0/0:5
 Up time: 00:12:57
 Input : 0 0 0 0
 Output: 0 0 0 0
 t1-2/0/0:6
 Up time: 00:12:57
 Input : 0 0 0 0
 Output: 0 0 0 0
Multilink detail statistics:
Bundle:
 Fragments:
 Input : 0 0 0 0
 Output: 0 0 0 0
 Non-fragments:
 Input : 0 0 0 0
 Output: 0 0 0 0
Protocol inet, MTU: 1500, Generation: 194, Route table: 0
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 11.1.1/24, Local: 11.1.1.2, Broadcast: Unspecified,
Generation: 151
DLCI 20
 Flags: Active, DCE-Configured
 Total down time: 00:03:18 sec, Last down: 00:15:38 ago
 Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
DLCI statistics:
 Active DLCI :2 Inactive DLCI :0

```

#### Related Documentation

- [Inline MLPPP for WAN Interfaces Overview on page 559](#)
- [Enabling MLPPP Link Fragmentation and Interleaving on page 769](#)
- [Example: Configuring Multilink Frame Relay FRF.16 on page 779](#)
- [Link and Multilink Services Interfaces Feature Guide for Routing Devices](#)
- [mlfr-uni-nni-bundles-inline on page 1412](#)
- [multi-link-layer-2-inline on page 1414](#)
- [pic \(MX Series Routers\)](#)
- [show interfaces \(Link Services IQ\) on page 1885](#)

# Configuring Additional Services on Link Services Interfaces

- [Configuring CoS on Link Services Interfaces on page 801](#)
- [Example: Configuring Link and Voice Services Interfaces with a Combination of Bundle Types on page 806](#)

## Configuring CoS on Link Services Interfaces

For link services IQ (**lsq-**) interfaces, Junos class of service (CoS) is fully supported and functions as described in the *Class of Service Feature Guide for Routing Devices*. For more information and detailed configuration examples, see [“Layer 2 Service Package Capabilities and Interfaces” on page 543](#).

On SRX Series devices, the **lsq-** interface is an internal interface, which is not associated with a physical interface. For information about link services on SRX Series devices, see the *Junos OS Interfaces Configuration Guide for Security Devices*.

For information about CoS functions and link services on M Series or T Series routers, see the following sections:

- [CoS for Link Services Interfaces on M Series and T Series Routers on page 801](#)
- [Example: Configuring CoS on Link Services Interfaces on page 803](#)

## CoS for Link Services Interfaces on M Series and T Series Routers

For Link Services PIC interfaces (**ls**) on M Series and T Series routers, queue 0 is the only queue that you should configure to receive fragmented packets. Configure all other queues to be higher-priority queues.

[Table 31 on page 801](#) summarizes how CoS queues work on link services (**ls**) interfaces.

Table 31: Link Services CoS Queues

| Supported Bundling Type   | Queue 0 | Higher-Priority Queues |
|---------------------------|---------|------------------------|
| Hash-based load balancing | No      | Yes                    |
| MLFR FRF.15               | Yes     | No                     |

Table 31: Link Services CoS Queues (*continued*)

| Supported Bundling Type | Queue 0 | Higher-Priority Queues |
|-------------------------|---------|------------------------|
| MLFR FRF.16             | Yes     | No                     |
| MLPPP                   | Yes     | No                     |

For M Series and T Series routers, CoS on link services (**ls**) interfaces works as follows:

- On all platforms, the Link Services PIC currently supports up to four queues: 0, 1, 2, and 3.
- Queue 0 uses MLFR FRF.15, MLFR FRF.16, or MLPPP to bundle packets.
- Higher-priority queues (1, 2, and 3) use hash-based load balancing to bundle packets. IP and MPLS header information is included in the hash.
- MLPPP packets traversing link services interfaces using queue 0 are fragmented and distributed across the constituent links. Queue 0 packets are sent on the least utilized link, proportional to its bandwidth. The queue 0 load balancer attempts to maintain even distribution of all traffic across all constituent links. In situations with a small number of high-priority traffic flows (queues 1, 2, and 3), queue 0 traffic might be unevenly distributed.
- For the MLFR FRF.16 protocol, only queue 0 works. If you configure a bundled interface to use MLFR FRF.16 with queue 0, then you must ensure the classifier does not send any traffic to queues 1, 2, and 3 on that interface.
- To carry high-priority traffic correctly on MLFR FRF.16 interfaces, you must configure an output firewall filter that forces all traffic into queue 0 on the **ls-fpc/pic/port.channel** interface.
- MLFR FRF.15 and MLPPP interfaces support CoS through packet interleaving. The MLFR FRF.16 standard does not support packet interleaving, so all packets destined for an FRF.16 PVC interface must egress from the same queue.
- For constituent link interfaces of Link Services PICs, you can configure standard scheduler maps.
- For input packets and fragments received from constituent links, you can use regular input firewall filters and standard CoS classifiers on the link services interface.
- For packets that pass through a link services interface and are destined for a constituent link interface, all traffic using queue 0 is fragmented. Traffic using higher-priority queues (1, 2, and 3) is not fragmented.
- For MLFR FRF.15 and MLPPP, routing protocol packets smaller than 128 bytes are sent to queue 3; routing protocol packets that exceed 128 bytes are sent to queue 0 and fragmented accordingly. For MLFR FRF.16, queue 0 is used for all packet sizes.
- You must configure output firewall classification for egress traffic on the link services interface, not directly on the constituent link interface directly.
- Inverse multiplexing for ATM (IMA) is not supported on link services interfaces.

For more information, see [“Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces” on page 773](#) and the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

## Example: Configuring CoS on Link Services Interfaces

Configure CoS on a link services interface and its constituent link interfaces.



**NOTE:** This example applies to M Series and T Series routers. For examples that apply to SRX Series devices, see the *Junos OS Interfaces Configuration Guide for Security Devices*.

Packets that do not match the firewall filters are sent to a queue that performs load balancing by sending fragments to all constituent links.

Packets that match the firewall filters are sent to a queue that does not support packet fragmentation and reassembly; instead, this traffic is load-balanced by sending each packet flow to a different constituent link. Each packet that matches a firewall filter is subjected to a hash on the IP source address and the IP destination address to determine the packet flow to which each packet belongs.

When you configure the MLPPP encapsulation type or the multilink FRF.15 Frame Relay end-to-end encapsulation type, routing protocol packets smaller than 128 bytes are sent to the network-control queue on the constituent link interface. This keeps routing protocols operating normally, even when low-speed links are congested by regular packets.

```
[edit interfaces]
ls-7/0/0 {
 unit 0 {
 encapsulation multilink-ppp;
 interleave-fragments;
 family inet {
 filter {
 output lfi_ls_filter;
 }
 address 10.54.0.2/32 {
 destination 10.54.0.1;
 }
 }
 }
}
ge-7/2/0 {
 unit 0 {
 family inet {
 address 192.168.1.1/24;
 }
 }
}
cel-7/3/6 {
 no-partition interface-type e1;
}
e1-7/3/6 {
```

```
encapsulation ppp;
unit 0 {
 family mppp {
 bundle ls-7/0/0.0;
 }
}
cel-7/3/7 {
 no-partition interface-type e1;
}
e1-7/3/7 {
 encapsulation ppp;
 unit 0 {
 family mppp {
 bundle ls-7/0/0.0;
 }
 }
}
[edit class-of-service]
classifiers {
 dscp dscp_default {
 import default;
 }
 inet-precedence inet-precedence_default {
 import default;
 }
}
code-point-aliases {
 dscp {
 af11 001010;
 af12 001100;
 af13 001110;
 af21 010010;
 af22 010100;
 af23 010110;
 af31 011010;
 af32 011100;
 af33 011110;
 af41 100010;
 af42 100100;
 af43 100110;
 be 000000;
 cs1 001000;
 cs2 010000;
 cs3 011000;
 cs4 100000;
 cs5 101000;
 cs6 110000;
 cs7 111000;
 ef 101110;
 }
 inet-precedence {
 af11 001;
 af21 010;
 af31 011;
 af41 100;
```

```
 be 000;
 cs6 110;
 cs7 111;
 ef 101;
 nc1 110;
 nc2 111;
 }
}
forwarding-classes {
 queue 0 be;
 queue 1 ef;
 queue 2 af;
 queue 3 nc;
}
interfaces {
 ge-7/2/0 {
 scheduler-map sched-map;
 unit 0 {
 classifiers {
 dscp dscp_default;
 }
 }
 }
 e1-7/3/6 {
 scheduler-map sched-map;
 }
 e1-7/3/7 {
 scheduler-map sched-map;
 }
 ls-7/0/0 {
 scheduler-map sched-map;
 unit 0 {
 classifiers {
 inet-precedence inet-precedence_default;
 }
 }
 }
}
scheduler-maps {
 sched-map {
 forwarding-class af scheduler af-scheduler;
 forwarding-class be scheduler be-scheduler;
 forwarding-class ef scheduler ef-scheduler;
 forwarding-class nc scheduler nc-scheduler;
 }
}
schedulers {
 af-scheduler {
 transmit-rate percent 25;
 buffer-size percent 25;
 }
 be-scheduler {
 transmit-rate percent 25;
 buffer-size percent 25;
 }
 ef-scheduler {
```

```
 transmit-rate percent 25;
 buffer-size percent 25;
 }
 nc-scheduler {
 transmit-rate percent 25;
 buffer-size percent 25;
 }
}
[edit firewall]
filter lfi_ls_filter {
 term term0 {
 from {
 destination-address {
 192.168.1.3/32;
 }
 precedence 5;
 }
 then {
 count count-192-168-1-3;
 forwarding-class af;
 accept;
 }
 }
 term default {
 then {
 log;
 forwarding-class best effort;
 accept;
 }
 }
}
```

- Related Documentation**
- [Link and Multilink Services Overview on page 711](#)
  - [Configuring Link Services Physical Interfaces on page 725](#)

## Example: Configuring Link and Voice Services Interfaces with a Combination of Bundle Types

---

```
[edit chassis]
fpc 1 {
 pic 3 {
 mlfr-uni-nni-bundles 4;
 }
}
[edit interfaces]
t1-0/2/0:0 {
 encapsulation multilink-frame-relay-uni-nni;
 unit 0 {
 family mlfr-uni-nni {
 bundle ls-1/3/0:0;
 }
 }
}
```



```
t1-0/2/0:1 {
 encapsulation multilink-frame-relay-uni-nni;
 unit 0 {
 family mlfr-uni-nni {
 bundle ls-1/3/0:0;
 }
 }
}
t1-0/2/0:5 {
 unit 0 {
 family mlppp {
 bundle ls-1/3/0:2;
 }
 }
}
t1-0/2/0:6 {
 unit 0 {
 family mlppp {
 bundle ls-1/3/0:2;
 }
 }
}
t1-0/2/0:7 {
 encapsulation frame-relay;
 unit 0 {
 dlci 20;
 family mlfr-end-to-end {
 bundle ls-1/3/0:1;
 }
 }
}
t1-0/2/0:8 {
 encapsulation frame-relay;
 unit 0 {
 dlci 20;
 family mlfr-end-to-end {
 bundle ls-1/3/0:1;
 }
 }
}
t1-0/2/0:10 {
 no-keepalives;
 encapsulation ppp;
 unit 0 {
 family mlppp {
 bundle lsq-1/1/0:0;
 }
 }
}
t3-1/0/0 {
 no-keepalives;
 encapsulation ppp;
 unit 0 {
 family mlppp {
 bundle lsq-1/1/0:2;
 }
 }
}
```

```
 }
 }
 lsq-1/1/0 {
 unit 0 {
 encapsulation multilink-ppp;
 compression {
 rtp {
 f-max-period 100;
 queues [q1 q2];
 port minimum 2000 maximum 6000;
 }
 }
 family inet {
 address 10.5.5.5/24;
 }
 }
 unit 1 {
 encapsulation multilink-ppp;
 compression {
 rtp {
 port minimum 2000 maximum 6000;
 }
 }
 family inet {
 address 10.6.6.1/24;
 }
 }
 unit 2 {
 encapsulation multilink-ppp;
 compression {
 rtp {
 port minimum 2000 maximum 6000;
 }
 }
 family inet {
 address 10.9.9.1/24;
 }
 }
 }
 t1-1/2/0 {
 no-keepalives;
 unit 0 {
 family mlppp {
 bundle lsq-1/1/0.1;
 }
 }
 }
 ls-1/3/0 {
 unit 1 {
 encapsulation multilink-frame-relay-end-to-end;
 family inet {
 address 10.1.4.1/24;
 }
 }
 unit 2 {
 encapsulation multilink-ppp;
 }
 }
}
```

```

 family inet {
 address 10.7.4.1/24;
 }
 }
}
ls-1/3/0:0 {
 encapsulation multilink-frame-relay-uni-nni;
 mlfr-uni-nni-bundle-options {
 debug-flags 15;
 }
 unit 0 {
 dlci 20;
 family inet {
 address 10.5.4.1/24;
 }
 }
}
[edit routing-options]
static {
 route 10.12.12.0/24 next-hop 10.1.1.9;
}

```

On Router B:

```

[edit chassis]
fpc 1 {
 pic 3 {
 mlfr-uni-nni-bundles 4;
 }
}
[edit interfaces]
ge-0/0/0 {
 unit 0 {
 family inet {
 address 10.1.1.1/24;
 }
 }
}
so-0/1/1 {
 encapsulation ppp;
 unit 0 {
 family inet {
 address 10.7.7.7/24;
 }
 }
}
t1-0/2/0:0 {
 encapsulation multilink-frame-relay-uni-nni;
 unit 0 {
 family mlfr-uni-nni {
 bundle ls-1/3/0:0;
 }
 }
}
t1-0/2/0:1 {
 encapsulation multilink-frame-relay-uni-nni;
}

```

```
 unit 0 {
 family mlfr-uni-nni {
 bundle ls-1/3/0:0;
 }
 }
 }
 t1-0/2/0:5 {
 no-keepalives;
 unit 0 {
 family mlppp {
 bundle ls-1/3/0.2;
 }
 }
 }
 t1-0/2/0:6 {
 no-keepalives;
 unit 0 {
 family mlppp {
 bundle ls-1/3/0.2;
 }
 }
 }
 t1-0/2/0:7 {
 dce;
 encapsulation frame-relay;
 unit 0 {
 dlci 20;
 family mlfr-end-to-end {
 bundle ls-1/3/0.1;
 }
 }
 }
 t1-0/2/0:8 {
 dce;
 encapsulation frame-relay;
 unit 0 {
 dlci 20;
 family mlfr-end-to-end {
 bundle ls-1/3/0.1;
 }
 }
 }
 t1-0/2/0:10 {
 no-keepalives;
 encapsulation ppp;
 unit 0 {
 family mlppp {
 bundle lsq-1/1/0.0;
 }
 }
 }
 t3-0/3/0 {
 no-keepalives;
 encapsulation ppp;
 unit 0 {
 family mlppp {
```

```
 bundle lsq-1/1/0.2;
 }
}
ge-1/0/0 {
 unit 0 {
 family inet {
 address 10.2.2.1/24;
 }
 }
}
lsq-1/1/0 {
 unit 0 {
 compression {
 rtp {
 port minimum 2000 maximum 6000;
 }
 }
 family inet {
 address 10.5.5.1/24;
 }
 }
 unit 1 {
 encapsulation multilink-ppp;
 compression {
 rtp {
 port minimum 16384 maximum 20102;
 }
 }
 family inet {
 address 10.3.4.1/24;
 }
 }
 unit 2 {
 encapsulation multilink-ppp;
 compression {
 rtp {
 port minimum 2000 maximum 6000;
 }
 }
 family inet {
 address 10.9.9.9/24;
 }
 }
}
t1-1/2/2 {
 no-keepalives;
 unit 0 {
 family mlppp {
 bundle ls-1/3/0.1;
 }
 }
}
t1-1/2/3 {
 no-keepalives;
 unit 0 {
```

```
 family mlppp {
 bundle lsq-1/1/0.1;
 }
 }
ls-1/3/0 {
 unit 1 {
 encapsulation multilink-frame-relay-end-to-end;
 family inet {
 address 10.1.4.4/24;
 }
 family iso;
 }
 unit 2 {
 encapsulation multilink-ppp;
 family inet {
 address 10.7.4.4/24;
 }
 }
}
ls-1/3/0:0 {
 dce;
 encapsulation multilink-frame-relay-uni-nni;
 unit 0 {
 dlci 20;
 family inet {
 address 10.5.4.4/24;
 }
 }
}
[edit routing-options]
static {
 route 10.12.12.0/24 next-hop 10.3.4.4;
}
```

**Related  
Documentation**

- [Configuring Link Services Physical Interfaces on page 725](#)
- [encapsulation \(Physical Interface\) on page 1604](#)

## PART 14

# Flow Monitoring and Flow Collection Services

- [Monitoring Traffic Using Active Flow Monitoring on page 815](#)
- [Monitoring Traffic Using Passive Flow Monitoring on page 829](#)
- [Processing and Exporting Multiple Records Using Flow Collection on page 839](#)





# Monitoring Traffic Using Active Flow Monitoring

- [Active Flow Monitoring Overview on page 815](#)
- [Configuring Flow Monitoring on page 818](#)
- [Example: Configuring Active Monitoring on Logical Systems on page 823](#)
- [Configuring Services Interface Redundancy with Flow Monitoring on page 826](#)
- [Flow Offloading on page 827](#)

## Active Flow Monitoring Overview

---

Using a Juniper Networks M Series Multiservice Edge or T Series Core Router or EX9200 switch, a selection of PICs (including the Monitoring Services PIC, Adaptive Services [AS] PIC, Multiservices PIC, or Multiservices DPC) and other networking hardware, you can monitor traffic flow and export the monitored traffic. Monitoring traffic allows you to do the following:

- Gather and export detailed information about IP version 4 (IPv4) traffic flows between source and destination nodes in your network.
- Sample all incoming IPv4 traffic on the monitoring interface and present the data in cflowd record format.
- Perform discard accounting on an incoming traffic flow.
- Encrypt or tunnel outgoing cflowd records, intercepted IPv4 traffic, or both.
- Direct filtered traffic to different packet analyzers and present the data in its original format (port mirror).



**NOTE:** Monitoring Services PICs, AS PICs, and Multiservices PICs must be mounted on an Enhanced Flexible PIC Concentrator (FPC) in an M Series or T Series router.

Multiservices DPCs installed in Juniper Networks MX Series 3D Universal Edge Routers support the same functionality, with the exception of the passive monitoring and flow-tap features.

Although the Monitoring Services PIC was designed initially for use as an offline passive flow monitoring tool, it can also be used in an active flow monitoring topology. In contrast, the AS or Multiservices PIC is designed exclusively for active flow monitoring. To use either the Monitoring Services PIC, AS PIC, or Multiservices PIC for active flow monitoring, you must install the PIC in an M Series or T Series router. The router participates in both the monitoring application and in the normal routing functionality of the network.

Starting with Junos OS Release 11.4, support for active monitoring is extended to logical systems running on T Series and MX Series routers. A logical system is a partition created from a physical router that performs independent routing tasks. Several logical systems in a single router with their own interfaces, policies, instances, and routing tables can perform functions handled by several different routers. A shared services PIC handles flows from all the logical systems. Only version 9 flows, IPv4, and MPLS templates are supported. See [“Example: Configuring Active Monitoring on Logical Systems” on page 823](#) for a sample configuration that enables active monitoring on a logical system.

Specified packets can be filtered and sent to the monitoring interface. For the Monitoring Services PIC, the interface name contains the **mo-** prefix. For the AS or Multiservices PIC, the interface name contains the **sp-** prefix.



**NOTE:** If you upgrade from the Monitoring Services PIC to the Adaptive Services or Multiservices PIC for active flow monitoring, you must change the name of your monitoring interface from **mo-fpc/pic/port** to **sp-fpc/pic/port**.

The major active flow monitoring actions you can configure at the **[edit forwarding-options]** hierarchy level are as follows:

- Sampling, with the **[edit forwarding-options sampling]** hierarchy. This option sends a copy of the traffic stream to an AS or Monitoring Services PIC, which extracts limited information (such as the source and destination IP address) from some of the packets in a flow. The original packets are forwarded to the intended destination as usual.
- Discard accounting, with the **[edit forwarding-options accounting]** hierarchy. This option quarantines unwanted packets, creates cflowd records that describe the packets, and discards the packets instead of forwarding them.
- Port mirroring, with the **[edit forwarding-options port-mirroring]** hierarchy. This option makes one full copy of all packets in a flow and delivers the copy to a single destination. The original packets are forwarded to the intended destination.
- Multiple port mirroring, with the **[edit forwarding-options next-hop-group]** hierarchy. This option allows multiple copies of selected traffic to be delivered to multiple destinations. (Multiple port mirroring requires a Tunnel Services PIC.)

Unlike passive flow monitoring, you do not need to configure a monitoring group. Instead, you can send filtered packets to a monitoring services or adaptive services interface (**mo-** or **sp-**) by using sampling or discard accounting. Optionally, you can configure port mirroring or multiple port mirroring to direct packets to additional interfaces.

These active flow monitoring options provide a wide variety of actions that can be performed on network traffic flows. However, the following restrictions apply:

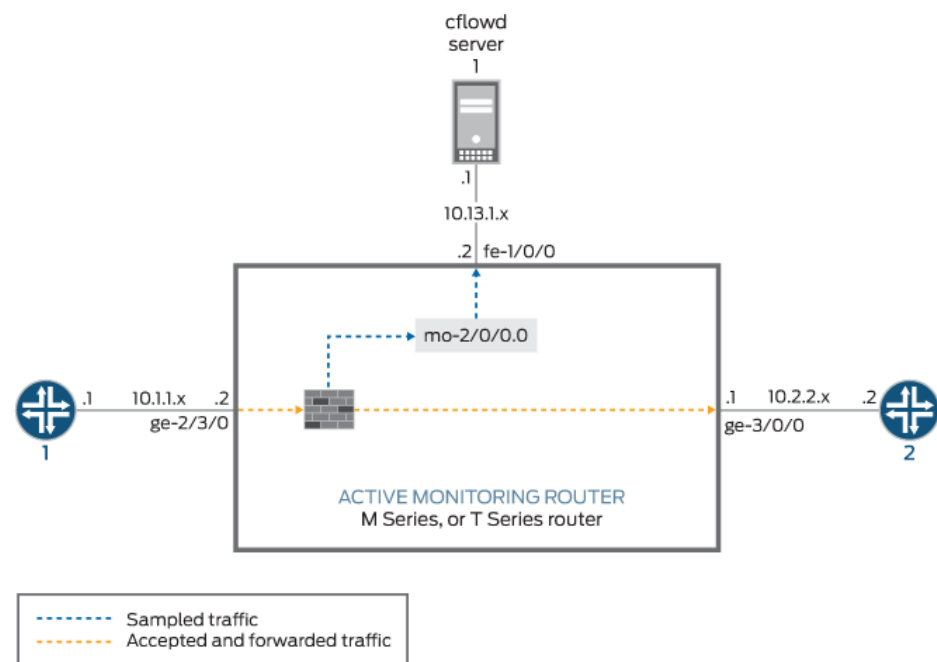
- The router or switch can perform sampling *or* port mirroring at any one time.
- The router or switch can perform forwarding *or* discard accounting at any one time.

Because the Monitoring Services, AS, and Multiservices PICs allow only one action to be performed at any one time, the following configuration options are available:

- Sampling and forwarding
- Sampling and discard accounting
- Port mirroring and forwarding
- Port mirroring and discard accounting
- Sampling and port mirroring on different sets of traffic

Figure 30 on page 817 shows a sample topology.

Figure 30: Active Monitoring Configuration Topology



8043214

In [Unresolved xref], traffic from Router 1 arrives on the monitoring router's Gigabit Ethernet **ge-2/3/0** interface. The exit interface on the monitoring router leading to destination Router 2 is **ge-3/0/0**, but this could be any interface type (such as SONET, Gigabit Ethernet, and so on). The export interface leading to the cflowd server is **fe-1/0/0**.

To enable active monitoring, configure a firewall filter on the interface **ge-2/3/0** with the following match conditions:

- Traffic matching certain firewall conditions is sent to the Monitoring Services PIC using filter-based forwarding. This traffic is quarantined and not forwarded to other routers.
- All other traffic is port-mirrored to the Monitoring Services PIC. Port mirroring copies each packet and sends the copies to the port-mirroring next hop (in this case, a Monitoring Services PIC). The original packets are forwarded out of the router as usual.

**Related  
Documentation**

- [Configuring Flow Monitoring on page 818](#)
- [Directing Replicated Flows to Multiple Flow Servers on page 926](#)
- [Configuring Services Interface Redundancy with Flow Monitoring on page 826](#)
- [Example: Configuring Active Monitoring on Logical Systems on page 823](#)

---

## Configuring Flow Monitoring

The flow-monitoring application performs traffic flow monitoring and enables lawful interception of traffic between two routers or switches. Traffic flows can either be passively monitored by an offline router or switch or actively monitored by a router participating in the network.

To configure flow monitoring you need to do the following:

- [Configuring Flow-Monitoring Interfaces on page 818](#)
- [Configuring Flow-Monitoring Properties on page 820](#)
- [Example: Configuring Flow Monitoring on page 822](#)

### Configuring Flow-Monitoring Interfaces

To enable flow monitoring on the Monitoring Services PIC, include the **mo-fpc/pic/port** statement at the **[edit interfaces]** hierarchy level:

```
mo-fpc/pic/port {
 unit logical-unit-number {
 family inet {
 address address {
 destination address;
 }
 filter {
 group filter-group-number;
 input filter-name;
 output filter-name;
 }
 sampling {
 [input output];
 }
 }
 }
 multiservice-options {
 (core-dump | no-core-dump);
 (syslog | no-syslog);
 flow-control-options {
```

```

 down-on-flow-control;
 dump-on-flow-control;
 reset-on-flow-control;
 }
}

```

Specify the physical and logical location of the flow-monitoring interface. You cannot use **unit 0**, because it is already used by internal processes. Specify the source and destination addresses. The **filter** statement allows you to associate an input or output filter or a filter group that you have already configured for this purpose. The **sampling** statement specifies the traffic direction: **input**, **output**, or both.

The **multiservice-options** statement allows you to configure properties related to flow-monitoring interfaces:

- Include the **core-dump** statement to enable storage of core files in **/var/tmp**.
- Include the **syslog** statement to enable storage of system logging information in **/var/log**.



**NOTE:** Boot images for monitoring services interfaces are specified at the [edit chassis images pic] hierarchy level. You must include the following configuration to make the flow monitoring feature operable:

```

[edit system]
ntp {
 boot-server ntp.juniper.net;
 server 172.17.28.5;
}
processes {
 ntp enable;
}

```

For more information, see the *Junos OS Administration Library for Routing Devices*.

- Include the **flow-control-options** statement to configure flow control.



**NOTE:** Starting with Junos OS Release 15.1, instead of an eJunos kernel core file, the multiservices PIC management daemon core file is generated when a prolonged flow control failure occurs and when you configure the setting to generate a core dump during prolonged flow control (by using the **dump-on-flow-control** option with the **flow-control-options** statement). The watchdog functionality continues to generate a kernel core file in such scenarios.

## Configuring Flow-Monitoring Properties

To configure flow-monitoring properties, include the **monitoring** statement at the **[edit forwarding-options]** hierarchy level:

```
monitoring name {
 family inet {
 output {
 cflowd hostname port port-number;
 export-format format;
 flow-active-timeout seconds;
 flow-export-destination {
 collector-pic;
 }
 flow-inactive-timeout seconds;
 interface interface-name {
 engine-id number;
 engine-type number;
 input-interface-index number;
 output-interface-index number;
 source-address address;
 }
 }
 }
}
```

A monitoring instance is a named entity that specifies collector information under the **monitoring name** statement. The following sections describe the properties you can configure:

- [Directing Traffic to Flow-Monitoring Interfaces on page 820](#)
- [Exporting Flows on page 821](#)
- [Configuring Time Periods when Flow Monitoring is Active and Inactive on page 821](#)

---

### Directing Traffic to Flow-Monitoring Interfaces

To direct traffic to a flow-monitoring interface, include the **interface** statement at the **[edit forwarding-options monitoring name output]** hierarchy level. By default, the Junos OS automatically assigns values for the **engine-id** and **engine-type** statements:

- **engine-id**—Monitoring interface location.
- **engine-type**—Platform-specific monitoring interface type.

The **source-address** statement specifies the traffic source for transmission of cflowd information; you must configure it manually. If you provide a different **source-address** statement for each monitoring services output interface, you can track which interface processes a particular cflowd record.

By default, the **input-interface-index** value is the SNMP index of the input interface. You can override the default by including a specific value. The **input-interface-index** and **output-interface-index** values are exported in fields present in the cflowd version 5 flow format.

## Exporting Flows

To direct traffic to a flow collection interface, include the **flow-export-destination** statement. For more information about flow collection, see *Flow Collection*.

To configure the cflowd version number, include the **export-format** statement at the **[edit forwarding-options monitoring name output]** hierarchy level. By default, version 5 is used. Version 8 enables the router software to aggregate the flow information using broader criteria and reduce cflowd traffic. Version 8 aggregation is performed periodically (every few seconds) on active flows and when flows are allowed to expire. Because the aggregation is performed periodically, active timeout events are ignored.

For more information on cflowd properties, see “Enabling Flow Aggregation” on page 898.

## Configuring Time Periods when Flow Monitoring is Active and Inactive

To configure time periods for active flow monitoring and intervals of inactivity, include the **flow-active-timeout** and **flow-inactive-timeout** statements at the **[edit forwarding-options monitoring name output]** hierarchy level:

- The **flow-active-timeout** statement specifies the time interval between flow exports for active flows. If the interval between the time the last packet was received and the time the flow was last exported exceeds the configured value, the flow is exported.

This timer is needed to provide periodic updates when a flow has a long duration. The active timeout setting enables the router to retain the start time for the flow as a constant and send out periodic cflowd reports. This in turn allows the collector to register the start time and determine that a flow has survived for a duration longer than the configured active timeout.



**NOTE:** In active flow monitoring, the cflowd records are exported after a time period that is a multiple of 60 seconds and greater than or equal to the configured active timeout value. For example, if the active timeout value is 90 seconds, the cflowd records are exported at 120-second intervals. If the active timeout value is 150 seconds, the cflowd records are exported at 180-second intervals, and so forth.

- The **flow-inactive-timeout** statement specifies the interval of inactivity for a flow that triggers the flow export. If the interval between the current time and the time that the last packet for this flow was received exceeds the configured inactive timeout value, the flow is allowed to expire.

If the flow stops transmitting for longer than the configured inactive timeout value, the router or switch purges it from the flow table and exports the cflowd record. As a result, the flow is forgotten as far as the PIC is concerned and if the same 5-tuple appears again, it is assigned a new start time and considered a new flow.

Both timers are necessary. The active timeout setting is needed to provide information for flows that constantly transmit packets for a long duration. The inactive timeout setting

enables the router or switch to purge flows that have become inactive and would waste tracking resources.



**NOTE:** The router must contain an Adaptive Services, Multiservices, or Monitoring Services PIC for the `flow-active-timeout` and `flow-inactive-timeout` statements to take effect.

---

## Example: Configuring Flow Monitoring

The following is an example of flow-monitoring properties configured to support input SONET/SDH interfaces, output monitoring services interfaces, and export to cflowd for flow analysis. To complete the configuration, you also need to configure the interfaces and set up a virtual private network (VPN) routing and forwarding (VRF) instance. For a complete example, see the *Junos OS, Release 14.2*. For information on cflowd, see [“Enabling Flow Aggregation” on page 898](#).

```
[edit forwarding-options]
monitoring group1 {
 family inet {
 output {
 cflowd 192.168.245.2 port 2055;
 export-format cflowd-version-5;
 flow-active-timeout 60;
 flow-inactive-timeout 30;
 interface mo-4/0/0.1 {
 engine-id 1;
 engine-type 1;
 input-interface-index 44;
 output-interface-index 54;
 source-address 192.168.245.1;
 }
 interface mo-4/1/0.1 {
 engine-id 2;
 engine-type 1;
 input-interface-index 45;
 output-interface-index 55;
 source-address 192.168.245.1;
 }
 interface mo-4/2/0.1 {
 engine-id 3;
 engine-type 1;
 input-interface-index 46;
 output-interface-index 56;
 source-address 192.168.245.1;
 }
 interface mo-4/3/0.1 {
 engine-id 4;
 engine-type 1;
 input-interface-index 47;
 output-interface-index 57;
 source-address 192.168.245.1;
 }
 }
 }
}
```



```

 }
 }
}

```

**Related Documentation**

- [Active Flow Monitoring Overview on page 815](#)
- [Directing Replicated Flows to Multiple Flow Servers on page 926](#)
- [Configuring Services Interface Redundancy with Flow Monitoring on page 826](#)
- [Example: Configuring Active Monitoring on Logical Systems on page 823](#)

## Example: Configuring Active Monitoring on Logical Systems

This example shows a sample configuration that allows you to configure active monitoring on a logical system. The following section shows the configuration on the master router:

```

[edit forwarding-options]
sampling {
 instance inst1 {
 input {
 rate 1;
 }
 family inet;
 output {
 flow-server 2.2.2.2 {
 port 2055;
 version9 {
 template {
 ipv4;
 }
 }
 }
 }
 }
 interface sp-0/1/0 {
 source-address 10.11.12.13;
 }
}
family mpls;
output {
 flow-server 2.2.2.2 {
 port 2055;
 version9 {
 template {
 mpls;
 }
 }
 }
 interface sp-0/1/0 {
 source-address 10.11.12.13;
 }
}
}

```

```
services {
 flow-monitoring {
 version9 {
 template ipv4 {
 flow-active-timeout 60;
 flow-inactive-timeout 60;
 ipv4-template;
 template-refresh-rate {
 packets 1000;
 seconds 10;
 }
 option-refresh-rate {
 packets 1000;
 seconds 10;
 }
 }
 }
 template mpls {
 mpls-template;
 }
 }
}
```

The configuration for the logical router uses the input parameters and the output interface for sampling from the master router. Each logical router should have separate template definitions for the flow-server configuration. The following section shows the configuration on the logical router:

```
logical-systems {
 ls-1 {
 firewall {
 family inet {
 filter test-sample {
 term term-1 {
 then {
 sample;
 accept;
 }
 }
 }
 }
 }
 interfaces {
 ge-0/0/1 {
 unit 0 {
 family inet {
 filter {
 input test-sample;
 output test-sample;
 }
 }
 }
 }
 }
 forwarding-options {
 sampling {
```

```

instance sample-inst1 {
 family inet;
 output {
 flow-server 2.2.2.2 {
 port 2055;
 version9 {
 template {
 ipv4-ls1;
 }
 }
 }
 }
}
family mpls;
output {
 flow-server 2.2.2.2 {
 port 2055;
 version9 {
 template {
 mpls-ls1;
 }
 }
 }
}
}
services {
 flow-monitoring {
 version9 {
 template ipv4-ls1 {
 flow-active-timeout 60;
 flow-inactive-timeout 60;
 ipv4-template;
 template-refresh-rate {
 packets 1000;
 seconds 10;
 }
 option-refresh-rate {
 packets 1000;
 seconds 10;
 }
 }
 template mpls-ls1 {
 mpls-template;
 }
 }
 }
}
}

```

- Related Documentation**
- [Active Flow Monitoring Overview on page 815](#)
  - [Configuring Flow Monitoring on page 818](#)

- [Directing Replicated Flows to Multiple Flow Servers on page 926](#)
- [Configuring Services Interface Redundancy with Flow Monitoring on page 826](#)

## Configuring Services Interface Redundancy with Flow Monitoring

Active monitoring services configurations on AS, Multiservices PICs, and Multiservices DPCs support redundancy. To configure redundancy, you specify a redundancy services PIC (**rsp**) interface in which the primary AS or Multiservices PIC is active and a secondary PIC is on standby. If the primary PIC fails, the secondary PIC becomes active, and all service processing is transferred to it. If the primary PIC is restored, it remains on standby and does not preempt the secondary PIC; you need to manually restore the services to the primary PIC. To determine which PIC is currently active, issue the **show interfaces redundancy** command.



**NOTE:** On flow-monitoring configurations, the only service option supported is *warm standby*, in which one backup PIC supports multiple working PICs. Recovery times are not guaranteed, because the configuration must be completely restored on the backup PIC after a failure is detected. However, configuration is preserved and available on the new active PIC.

As with the other services that support warm standby, you can issue the **request interfaces (revert | switchover)** command to switch manually between the primary and secondary flow monitoring interfaces.

For more information, see “[Configuring AS or Multiservices PIC Redundancy](#)” on page 41. For information on operational mode commands, see the [CLI Explorer](#).

A sample configuration follows.

```
interface {
 rsp0 {
 redundancy-options {
 primary sp-0/0/0;
 secondary sp-1/3/0;
 }
 unit 0 {
 family inet;
 }
 }
}
interface {
 ge-0/2/0 {
 unit 0 {
 family inet {
 filter {
 input as_sample;
 }
 }
 }
 address 10.58.255.49/28;
```

```

 }
 }
}
forwarding-options {
 sampling {
 instance instance1 { # named instances of sampling parameters
 input {
 rate 1;
 run-length 0;
 max-packets-per-second 65535;
 }
 family inet {
 output {
 flow-server 10.10.10.2 {
 port 5000;
 version 5;
 }
 flow-active-timeout 60;
 interface rsp0 {
 source-address 10.10.10.1;
 }
 }
 }
 }
 }
}
}
firewall {
 filter as_sample {
 term t1 {
 then {
 sample;
 accept;
 }
 }
 }
}
}

```

#### Related Documentation

- [Active Flow Monitoring Overview on page 815](#)
- [Configuring Flow Monitoring on page 818](#)
- [Directing Replicated Flows to Multiple Flow Servers on page 926](#)
- [Example: Configuring Active Monitoring on Logical Systems on page 823](#)

## Flow Offloading

The Junos OS enables you to configure flow offloading for PICS on MX Series routers using Modular Port Concentrator (MPCs) with Modular Interface Cards (MICs). Flows are offloaded to Fast Update Filters (FUFs) on the Packet Forwarding Engine. Offloading produces the greatest benefits when applied to long-lasting or high-bandwidth flows.

The maximum number of active offloads is 200,000 per PIC. When offloaded flows are deleted, more flows can be offloaded.

To configure flow offloading:

- At the **[edit interfaces *interface-name* services-options]** hierarchy level, enter the **trio-flow-offload minimum-bytes *minimum-bytes*** statement.

```
user@host# edit services interface-name
[edit services interface-name services-options]
user@host# set trio-flow-offload minimum-bytes minimum-bytes
```

In the following example, flows are offloaded when they consist of no less than 1024 bytes:

```
user@host# edit services ms-0/1/0
[edit services ms-0/1/0 services-options]
user@host# set trio-flow-offload minimum-bytes 1024
```

**Related Documentation**

- [trio-flow-offload on page 1777](#)

# Monitoring Traffic Using Passive Flow Monitoring

- [Passive Flow Monitoring Overview on page 829](#)
- [Enabling Passive Flow Monitoring on page 830](#)

## Passive Flow Monitoring Overview

---

Using a Juniper Networks M Series Multiservice Edge or T Series Core Router, a selection of PICs (including the Monitoring Services PIC, Adaptive Services [AS] PIC, Multiservices PIC, or Multiservices DPC) and other networking hardware, you can monitor traffic flow and export the monitored traffic. Monitoring traffic allows you to do the following:

- Gather and export detailed information about IP version 4 (IPv4) traffic flows between source and destination nodes in your network.
- Sample all incoming IPv4 traffic on the monitoring interface and present the data in cflowd record format.
- Perform discard accounting on an incoming traffic flow.
- Encrypt or tunnel outgoing cflowd records, intercepted IPv4 traffic, or both.
- Direct filtered traffic to different packet analyzers and present the data in its original format (port mirror).




**NOTE:** Monitoring Services PICs, AS PICs, and Multiservices PICs must be mounted on an Enhanced Flexible PIC Concentrator (FPC) in an M Series or T Series router.

Multiservices DPCs installed in Juniper Networks MX Series 3D Universal Edge Routers support the same functionality, with the exception of the passive monitoring and flow-tap features.

The router used for passive monitoring does not route packets from the monitored interface, nor does it run any routing protocols related to those interfaces; it only receives traffic flows, collects intercepted traffic, and exports it to cflowd servers and packet analyzers. [Figure 31 on page 830](#) shows a typical topology for the passive flow-monitoring application.

Figure 31: Passive Monitoring Application Topology



Traffic travels normally between Router 1 and Router 2. To redirect IPv4 traffic, you insert an optical splitter on the interface between these two routers. The optical splitter copies and redirects the traffic to the monitoring station, which is an M40e, M160, M320, or T Series router. The optical cable connects only the receive port on the monitoring station, never the transmit port. This configuration allows the monitoring station to receive traffic from the router being monitored but never to transmit it back.

If you are monitoring traffic flow, the Internet Processor II application-specific integrated circuit (ASIC) in the router forwards a copy of the traffic to the Monitoring Services, Adaptive Services, or Multiservices PIC in the monitoring station. If more than one monitoring PIC is installed, the monitoring station distributes the load of the incoming traffic across the multiple PICs. The monitoring PICs generate flow records in cflowd version 5 format, and the records are then exported to the cflowd collector.

If you are performing lawful interception of traffic between the two routers, the Internet Processor II ASIC filters the incoming traffic and forwards it to the Tunnel Services PIC. Filter-based forwarding is then applied to direct the traffic to the packet analyzers.

Optionally, the intercepted traffic or the cflowd records can be encrypted by the ES PIC or IP Security (IPsec) services and then sent to a cflowd server or packet analyzer.

**Related Documentation**

- [Enabling Passive Flow Monitoring on page 830](#)

## Enabling Passive Flow Monitoring

You can monitor IPv4 traffic from another router if you have the following components installed in an M Series, MX Series, or T Series router:

- Monitoring Services, Adaptive Services, or Multiservices PICs to perform the service processing
- SONET/SDH, Fast Ethernet, or Gigabit Ethernet PICs as transit interface



On SONET/SDH interfaces, you enable passive flow monitoring by including the **passive-monitor-mode** statement at the **[edit interfaces so-fpc/pic/port unit logical-unit-number]** hierarchy level:

```
[edit interfaces so-fpc/pic/port unit logical-unit-number]
passive-monitor-mode;
```

On Asynchronous Transfer Mode (ATM), Fast Ethernet, or Gigabit Ethernet interfaces, you enable passive flow monitoring by including the **passive-monitor-mode** statement at the **[edit interfaces interface-name]** hierarchy level:

```
[edit interfaces interface-name]
passive-monitor-mode;
```

IPv6 passive monitoring is not supported on Monitoring Services PICs. You must configure port mirroring to forward the packets from the passive monitored ports to other interfaces. Interfaces configured on the following FPCs and PIC support IPv6 passive monitoring on the T640 and T1600 routers:

- Enhanced Scaling FPC2
- Enhanced Scaling FPC3
- Enhanced II FPC1
- Enhanced II FPC2
- Enhanced II FPC3
- Enhanced Scaling FPC4
- Enhanced Scaling FPC4.1
- 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP (supported on both WAN-PHY and LAN-PHY mode for both IPv4 and IPv6 addresses)
- Gigabit Ethernet PIC with SFP
- 10-Gigabit Ethernet PIC with XENPAK (T1600 router)
- SONET/SDH OC192/STM64 PIC (T1600 router)
- SONET/SDH OC192/STM64 PICs with XFP (T1600 router)
- SONET/SDH OC48c/STM16 PIC with SFP (T1600 router)
- SONET/SDH OC48/STM16 (Multi-Rate)
- SONET/SDH OC12/STM4 (Multi-Rate) PIC with SFP
- Type 1 SONET/SDH OC3/STM1 (Multi-Rate) PIC with SFP

To configure port mirroring, include the **port-mirroring** statement at the **[edit forwarding-options]** hierarchy level.

When you configure an interface in passive monitoring mode, the Packet Forwarding Engine silently drops packets coming from that interface and destined to the router itself. Passive monitoring mode also stops the Routing Engine from transmitting any packet from that interface. Packets received from the monitored interface can be forwarded to monitoring interfaces. If you include the **passive-monitor-mode** statement in the configuration:

- The ATM interface is always up, and the interface does not receive or transmit incoming control packets, such as Operation, Administration, and Maintenance (OAM) and Interim Local Management Interface (ILMI) cells.
- The SONET/SDH interface does not send keepalives or alarms and does not participate actively on the network.
- Gigabit and Fast Ethernet interfaces can support both per-port passive monitoring and per-VLAN passive monitoring. The destination MAC filter on the receive port of the Ethernet interfaces is disabled.
- Ethernet encapsulation options are not allowed.
- Ethernet interfaces do not support the **stacked-vlan-tagging** statement for both IPv4 and IPv6 packets in passive monitoring mode.

On monitoring services interfaces, you enable passive flow monitoring by including the **family** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level, specifying the **inet** option:

```
[edit interfaces interface-name unit logical-unit-number]
family inet;
```

For the monitoring services interface, you can configure multiservice physical interface properties. For more information, see [“Configuring Flow-Monitoring Interfaces” on page 818](#).

For conformity with the cflowd record structure, you must include the **receive-options-packets** and **receive-ttl-exceeded** statements at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet]
receive-options-packets;
receive-ttl-exceeded;
```

For more information, see the following sections:

- [Passive Flow Monitoring for MPLS Encapsulated Packets on page 832](#)
- [Example: Enabling IPv4 Passive Flow Monitoring on page 834](#)
- [Example: Enabling IPv6 Passive Flow Monitoring on page 836](#)

## Passive Flow Monitoring for MPLS Encapsulated Packets

On monitoring services interfaces, you can process MPLS packets that have not been assigned label values and have no corresponding entry in the **mpls.0** routing table. This allows you to assign a default route to unlabeled MPLS packets.

To configure a default label value for MPLS packets, include the **default-route** statement at the **[edit protocols mpls interface *interface-name* label-map]** hierarchy level:

```
[edit protocols mpls interface interface-name label-map]
default-route {
 (next-hop (address | interface-name | address/interface-name)) | (reject | discard);
 (pop | (swap <out-label>));
 class-of-service value;
 preference preference;
 type type;
}
```

For more information about static labels, see the *MPLS Applications Feature Guide for Routing Devices*.

### Removing MPLS Labels from Incoming Packets

The Junos OS can forward only IPv4 packets to a Monitoring Services, Adaptive Services, or Multiservices PIC. IPv4 and IPv6 packets with MPLS labels cannot be forwarded to a monitoring PIC. By default, if packets with MPLS labels are forwarded to the monitoring PIC, they are discarded. To monitor IPv4 and IPv6 packets with MPLS labels, you must remove the MPLS labels as the packets arrive on the interface.

You can remove up to two MPLS labels from an incoming packet by including the **pop-all-labels** statement at the [edit interfaces *interface-name* (atm-options | fastether-options | gigether-options | sonet-options) mpls] hierarchy level:

```
[edit interfaces interface-name (atm-options | fastether-options | gigether-options |
sonet-options) mpls]
pop-all-labels {
 required-depth [numbers];
}
```

By default, the **pop-all-labels** statement takes effect for incoming packets with one or two labels. You can specify the number of MPLS labels that an incoming packet must have for the **pop-all-labels** statement to take effect by including the **required-depth** statement at the [edit interfaces *interface-name* (atm-options | fastether-options | gigether-options | sonet-options) mpls pop-all-labels] hierarchy level:

```
[edit interfaces interface-name (atm-options | fastether-options | gigether-options |
sonet-options) mpls pop-all-labels]
required-depth [numbers];
```

The required depth can be 1, 2, or [ 1 2 ]. If you include the **required-depth 1** statement, the **pop-all-labels** statement takes effect for incoming packets with one label only. If you include the **required-depth 2** statement, the **pop-all-labels** statement takes effect for incoming packets with two labels only. If you include the **required-depth [ 1 2 ]** statement, the **pop-all-labels** statement takes effect for incoming packets with one or two labels. A required depth of [ 1 2 ] is equivalent to the default behavior of the **pop-all-labels** statement.

When you remove MPLS labels from incoming packets, note the following:

- The **pop-all-labels** statement has no effect on IP packets with three or more MPLS labels.
- When you enable MPLS label removal, you must configure all ports on a PIC with the same label popping mode and required depth.

- You use the **pop-all-labels** statement to enable passive monitoring applications, not active monitoring applications.
- You cannot apply MPLS filters or accounting to the MPLS labels because the labels are removed as soon as the packet arrives on the interface.
- On ATM2 interfaces, you must use a label value greater than 4095 because the lower range of MPLS labels is reserved for label-switched interface (LSI) and virtual private LAN service (VPLS) support. For more information, see the *Junos OS VPNs Library for Routing Devices*.
- The following ATM encapsulation types are not supported on interfaces with MPLS label removal:
  - **atm-ccc-cell-relay**
  - **atm-ccc-vc-mux**
  - **atm-mlppp-llc**
  - **atm-tcc-snap**
  - **atm-tcc-vc-mux**
  - **ether-over-atm-llc**
  - **ether-vpls-over-atm-llc**

### Example: Enabling IPv4 Passive Flow Monitoring

The following example shows a complete configuration for enabling passive flow monitoring on an Ethernet interface.

In this example, the Gigabit Ethernet interface can accept all Ethernet packets. It strips VLAN tags (if there are any) and up to two MPLS labels blindly, and passes IPv4 packets to the monitoring interface. With this configuration, it can monitor IPv4, VLAN+IPv4, VLAN+MPLS+IPv4, and VLAN+MPLS+MPLS+IPv4 labeled packets.

The Fast Ethernet interface can accept only packets with VLAN ID 100. All other packets are dropped. With this configuration, it can monitor VLAN (ID=100)+IPv4, VLAN (ID=100)+MPLS+IPv4, and VLAN (ID=100)+MPLS+MPLS+IPv4 labeled packets.

```
[edit firewall]
family inet {
 filter input-monitoring-filter {
 term def {
 then {
 count counter;
 accept;
 }
 }
 }
}
[edit interfaces]
ge-0/0/0 {
 passive-monitor-mode;
 gigether-options {
```

```

 mpls {
 pop-all-labels;
 }
 }
 unit 0 {
 family inet {
 filter {
 input input-monitoring-filter;
 }
 }
 }
}
fe-0/1/0 {
 passive-monitor-mode;
 vlan-tagging;
 fastether-options {
 mpls {
 pop-all-labels required-depth [1 2];
 }
 }
 unit 0 {
 vlan-id 100;
 family inet {
 filter {
 input input-monitoring-filter;
 }
 }
 }
}
mo-1/0/0 {
 unit 0 {
 family inet {
 receive-options-packets;
 receive-ttl-exceeded;
 }
 }
 unit 1 {
 family inet;
 }
}
[edit forwarding-options]
monitoring mon1 {
 family inet {
 output {
 export-format cflowd-version-5;
 cflowd 50.0.0.2 port 2055;
 interface mo-1/0/0.0 {
 source-address 50.0.0.1;
 }
 }
 }
}
[edit routing-instances]
monitoring-vrf {
 instance-type vrf;
 interface ge-0/0/0.0;
}

```

```
interface fe-0/1/0.0;
interface mo-1/0/0.1;
route-distinguisher 68:1;
vrf-import monitoring-vrf-import;
vrf-export monitoring-vrf-export;
routing-options {
 static {
 route 0.0.0.0/0 next-hop mo-1/0/0.1;
 }
}
[edit policy-options]
policy-statement monitoring-vrf-import {
 then {
 reject;
 }
}
policy-statement monitoring-vrf-export {
 then {
 reject;
 }
}
```

### Example: Enabling IPv6 Passive Flow Monitoring

The following example shows a complete configuration for enabling IPv6 passive flow monitoring on an Ethernet interface.

In this example, the Gigabit Ethernet interface can accept all Ethernet packets. It strips VLAN tags (if there are any) and up to two MPLS labels blindly, and passes IPv6 packets to the monitoring interface. With this configuration, the Gigabit Ethernet interface can monitor IPv6, VLAN+IPv6, VLAN+MPLS+IPv6, and VLAN+MPLS+MPLS+IPv6 labeled packets.

The vlan-tagged Gigabit Ethernet interface can accept only packets with VLAN ID 100. All other packets are dropped. With this configuration, it can monitor VLAN (ID=100)+IPv6, VLAN (ID=100)+MPLS+IPv6, and VLAN (ID=100)+MPLS+MPLS+IPv6 labeled packets.

```
[edit interfaces]
xe-0/1/0 {
 passive-monitor-mode;
 unit 0 {
 family inet6 {
 filter {
 input port-mirror6;
 }
 address 2001::1/128;
 }
 }
}
xe-0/1/2 {
 passive-monitor-mode;
 vlan-tagging;
 unit 0 {
 vlan-id 100;
 }
}
```

```

 family inet6 {
 filter {
 input port-mirror6;
 }
 }
 }
}
xe-0/1/1 {
 unit 0 {
 family inet6 {
 address 2000::1/128;
 }
 }
}
[edit firewall]
family inet6 {
 filter port-mirror6 {
 term term2 {
 then {
 count count_pm;
 port-mirror;
 accept;
 }
 }
 }
}
[edit forwarding options]
port-mirroring {
 input {
 rate 1;
 }
 family inet6 {
 output {
 interface xe-0/1/1.0 {
 next-hop 2000::3;
 }
 no-filter-check;
 }
 }
}
}

```

**Related Documentation**

- [Passive Flow Monitoring Overview on page 829](#)





# Processing and Exporting Multiple Records Using Flow Collection

- [Flow Collection Overview on page 839](#)
- [Configuring Flow Collection on page 840](#)
- [Sending cflowd Records to Flow Collector Interfaces on page 843](#)
- [Configuring Flow Collection Mode and Interfaces on Services PICs on page 844](#)

## Flow Collection Overview

---

You can process and export multiple cflowd records with a flow collector interface. You create a flow collector interface on a Monitoring Services II or Multiservices 400 PIC. The flow collector interface combines multiple cflowd records into a compressed ASCII data file and exports the file to an FTP server. To convert a services PIC into a flow collector interface, include the **flow-collector** statement at the **[edit chassis fpc fpc-slot pic pic-slot monitoring-services application]** hierarchy level.

You can use the services PIC for either flow collection or monitoring, but not for both types of service simultaneously. When converting the PIC between service types, you must configure the **flow-collector** statement, take the PIC offline, and then bring the PIC back online. Restarting the router does not enable the new service type.

A flow collector interface, designated by the **cp-fpc/pic/port** interface name, requires three logical interfaces for correct operation. Units 0 and 1 are used to send the compressed ASCII data files to an FTP server, while Unit 2 is used to receive cflowd records from a monitoring services interface.



**NOTE:** Unlike conventional interfaces, the **address** statement at the **[edit interfaces cp-fpc/pic/port unit unit-number family inet]** hierarchy level corresponds to the IP address of the Routing Engine. Likewise, the **destination** statement at the **[edit interfaces cp-fpc/pic/port unit unit-number family inet address ip-address]** hierarchy level corresponds to the IP address of the flow collector interface. As a result, you must configure the **destination** statement for Unit 0 and 1 with *local* addresses that can reach the FTP server. Similarly, configure the **destination** statement for Unit 2 with a *local* IP address so it can reach the monitoring services interface that sends cflowd records.

To activate flow collector services after the services PIC is converted into a flow collector, include the **flow-collector** statement at the **[edit services]** hierarchy level.

After you activate the flow collector, you need to configure the following components:

- Destination of the FTP server
- File specifications
- Input interface-to-flow collector interface mappings
- Transfer log settings

**Related  
Documentation**

- [Configuring Flow Collection on page 840](#)
- [Sending cflowd Records to Flow Collector Interfaces on page 843](#)
- [Configuring Flow Collection Mode and Interfaces on Services PICs on page 844](#)

---

## Configuring Flow Collection

This section describes the following tasks for configuring flow collection:

- [Configuring Destination FTP Servers for Flow Records on page 840](#)
- [Configuring a Packet Analyzer on page 841](#)
- [Configuring File Formats on page 841](#)
- [Configuring Interface Mappings on page 842](#)
- [Configuring Transfer Logs on page 842](#)
- [Configuring Retry Attempts on page 843](#)

### Configuring Destination FTP Servers for Flow Records

Flow collection destinations are where the compressed ASCII data files are sent after the cflowd records are collected and processed. To specify the destination FTP server, include the **destinations** statement at the **[edit services flow-collector]** hierarchy level. You can specify up to two FTP server destinations and include the password for each configured server. If two FTP servers are configured, the first server in the configuration is the primary server and the second is a backup server.

To configure a destination for flow collection files, include the **destinations** statement at the **[edit services flow-collector]** hierarchy level:

```
[edit services flow-collector]
destinations {
 ftp:url {
 password "password";
 }
}
```

To specify the destination FTP server, include the **ftp:url** statement. The value **url** is the FTP server address for the primary flow collection destination and can include macros.

When you include macros in the **ftp:url** statement, a directory can be created only for a single level. For example, the path **ftp://10.2.2.2/%m/%Y** expands to **ftp://10.2.2.2/01/2005**, and the software attempts to create the directory **01/2005** on the destination FTP server. If the **01/** directory already exists on the destination FTP server, the software creates the **/2005/ directory** one level down. If the **01/** directory does not exist on the destination FTP server, the software cannot create the **/2005/ directory**, and the FTP server destination will fail. For more information about macros, see [ftp](#).

To specify the FTP server password, include the **password "password"** statement. The password must be enclosed in quotation marks. You can specify up to two destination FTP servers. The first destination specified is considered the primary destination.

## Configuring a Packet Analyzer

You can specify values for the IP address and identifier of a packet analyzer to which the flow collector interface sends traffic for analysis. The values you specify here override any default values configured elsewhere.

To configure an IP address and identifier for the packet analyzer, include the **analyzer-address** and **analyzer-id** statements at the **[edit services flow-collector]** hierarchy level:

```
[edit services flow-collector]
analyzer-address address;
analyzer-id name;
```

## Configuring File Formats

You configure data file formats, name formats, and transfer characteristics for the flow collection files. File records are sent to the destination FTP server when the timer expires or when a preset number of records are received, whichever comes first.

To configure the flow collection file format, include the **file-specification** statement at the **[edit services flow-collector]** hierarchy level:

```
[edit services flow-collector]
file-specification {
 variant variant-number {
 data-format format;
 name-format format;
 transfer {
 record-level number;
 timeout seconds;
 }
 }
}
```

To set the data file format, include the **data-format** statement. To set the file name format, include the **name-format** statement. To set the export timer and file size thresholds, include the **transfer** statement and specify values for the **timeout** and **record-level** options.

For example, you can specify the name format as follows:

```
[edit services flow-collector file-specification variant variant-number]
name-format "cFlowd-py69Ni69-0-%D_%T-%I_%N.bcp.bi.gz";
```

In this example, **cFlowd-py69Ni69-0** is the static portion used verbatim, **%D** is the date in YYYYMMDD format, **%T** is the time in HHMMSS format, **%I** is the value of **ifAlias**, **%N** is the generation number, and **bcp.bi.gz** is a user-configured string. A number of macros are supported for expressing the date and time information in different ways; for a complete list, see the summary section for [name-format](#).

## Configuring Interface Mappings

You can match an input interface with a flow collector interface and apply the preset file specifications to the input interface.

To configure an interface mapping, include the **interface-map** statement at the **[edit services flow-collector]** hierarchy level:

```
[edit services flow-collector]
interface-map {
 collector interface-name;
 file-specification variant-number;
 interface-name {
 collector interface-name;
 file-specification variant-number;
 }
}
```

To configure the default flow collector and file specifications for all input interfaces, include the **file-specification** and **collector** statements at the **[edit services flow-collector interface-map]** hierarchy level. To override the default settings and apply flow collector and file specifications to a specific input interface, include the **file-specification** and **collector** statements at the **[edit services flow-collector interface-map *interface-name*]** hierarchy level.

## Configuring Transfer Logs

You can configure the filename, export interval, maximum size, and destination FTP server for log files containing the transfer activity history for a flow collector interface.

To configure a transfer log, include the **transfer-log-archive** statement at the **[edit services flow-collector]** hierarchy level:

```
[edit services flow-collector]
transfer-log-archive {
 archive-sites {
 ftp:url {
 password "password";
 username username;
 }
 }
 filename-prefix prefix;
 maximum-age minutes;
}
```

To configure the destination for archiving files, include the **archive-sites** statement. Specify the filename as follows:

```
[edit services flow-collector transfer-log]
filename "cFlowd-py69Ni69-0-%D_%T";
```

where **cFlowd-py69Ni69-0** is the static portion used verbatim, **%D** is the date in YYYYMMDD format, and **%T** is the time in HHMMSS format.

You can optionally include the following statements:

- **filename-prefix**—Sets a standard prefix for all the logged files.
- **maximum-age**—Specifies the duration a file remains on the server. The range is 1 through 360 minutes.

## Configuring Retry Attempts

You can specify values for situations in which the flow collector interface needs more than one attempt to transfer log files to the FTP server:

- Maximum number of retry attempts
- Amount of time the flow collector interface waits between successive retries

To configure retry settings, include the **retry** and **retry-delay** statements at the **[edit services flow-collector]** hierarchy level:

```
retry number;
retry-delay seconds;
```

The **retry** value can be from 0 through 10. The **retry-delay** value can be from 0 through 60 seconds.

### Related Documentation

- [Flow Collection Overview on page 839](#)
- [Sending cflowd Records to Flow Collector Interfaces on page 843](#)
- [Configuring Flow Collection Mode and Interfaces on Services PICs on page 844](#)
- *Example: Configuring Flow Collection*

---

## Sending cflowd Records to Flow Collector Interfaces

To specify a flow collector interface as the destination for cflowd records coming from a services PIC, include the **collector-pic** statement at the **[edit forwarding-options monitoring group-name family inet output flow-export-destination]** hierarchy level:

```
[edit forwarding-options monitoring group-name family inet output flow-export-destination]
collector-pic;
```

You can select either the flow collector interface or a cflowd server as the destination for cflowd records, but not both at the same time.

- Related Documentation**
- [Flow Collection Overview on page 839](#)
  - [Configuring Flow Collection on page 840](#)
  - [Configuring Flow Collection Mode and Interfaces on Services PICs on page 844](#)
  - *Example: Configuring Flow Collection*

---

## Configuring Flow Collection Mode and Interfaces on Services PICs

---

You can select the services PIC to run in either flow collection mode or monitoring mode, but not both.

To set the services PIC to run in flow collection mode, include the **flow-collector** statement at the **[edit chassis fpc slot-number pic pic-number monitoring-services application]** hierarchy level:

```
[edit chassis fpc slot-number pic pic-number monitoring-services application]
flow-collector;
```

For further information on configuring chassis properties, see the *Junos OS Administration Library for Routing Devices*.

To specify flow collection interfaces, you configure the **cp** interface at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
cp-fpc/pic/port {
 ...
}
```

- Related Documentation**
- [Flow Collection Overview on page 839](#)
  - [Configuring Flow Collection on page 840](#)
  - [Sending cflowd Records to Flow Collector Interfaces on page 843](#)
  - *Example: Configuring Flow Collection*

## PART 15

# Flow Capture Services

- [Dynamically Capturing Packet Flows Using Junos Capture Vision on page 847](#)
- [Detecting Threats and Intercepting Flows Using Junos Packet Vision on page 859](#)





## CHAPTER 58

# Dynamically Capturing Packet Flows Using Junos Capture Vision

- [Understanding Junos Capture Vision on page 847](#)
- [Configuring Junos Capture Vision on page 849](#)
- [Example: Configuring Junos Capture Vision on page 855](#)

## Understanding Junos Capture Vision

---

Junos Capture Vision (known as dynamic flow capture in Junos OS Releases earlier than 13.2) enables you to capture packet flows on the basis of dynamic filtering criteria. Specifically, you can use this feature to forward passively monitored packet flows that match a particular filter list to one or more destinations using an on-demand control protocol.

This topic contains the following sections:

- [Junos Capture Vision Architecture on page 847](#)
- [Liberal Sequence Windowing on page 848](#)
- [Intercepting IPv6 Flows on page 849](#)

## Junos Capture Vision Architecture

The architecture consists of one or more *control sources* that send requests to a Juniper Networks router to monitor incoming data, and then forward any packets that match specific filter criteria to a set of one or more *content destinations*. The architectural components are defined as follows:

- **Control source**—A client that monitors electronic data or voice transfer over the network. The control source sends filter requests to the Juniper Networks router using the Dynamic Task Control Protocol (DTCP), specified in draft-cavuto-dtcp-03.txt at <http://www.ietf.org/internet-drafts>. The control source is identified by a unique identifier and an optional list of IP addresses.
- **Monitoring platform**—A T Series or M320 router containing one or more Dynamic Flow Capture (DFC) PICs, which support dynamic flow capture processing. The monitoring platform processes the requests from the control sources, creates the filters, monitors

incoming data flows, and sends the matched packets to the appropriate content destinations.

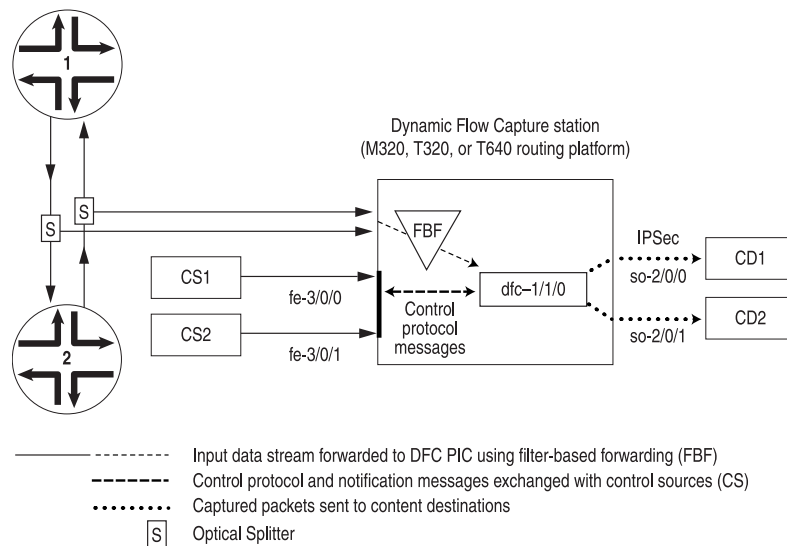
- Content destination—Recipient of the matched packets from the monitoring platform. Typically the matched packets are sent using an IP Security (IPsec) tunnel from the monitoring platform to another router connected to the content destination. The content destination and the control source can be physically located on the same host. For more information on IPsec tunnels, see *Junos VPN Site Secure*.



**NOTE:** The Junos Capture Vision PIC (either a Monitoring Services III PIC or Multiservices 400 PIC) forwards the entire packet content to the content destination, rather than to a content record as is done with cflowd or flow aggregation version 9 templates.

Figure 32 on page 848 shows a sample topology. The number of control sources and content destinations is arbitrary.

**Figure 32: Junos Capture Vision Topology**



g017075

## Liberal Sequence Windowing

Each DTCP packet (add, delete, list, and refresh packets) contains a 64-bit sequence number to identify the order of the packets. Because the network is connectionless, the DTCP packets can arrive out of order to the router running the Junos Capture Vision application.

The *liberal sequence window* feature implements a negative window for the sequence numbers received in the DTCP packets. It enables the Junos Capture Vision application to accept not only DTCP packets with sequence numbers greater than those previously received, but also DTCP packets with lesser sequence numbers, up to a certain limit. This limit is the negative window size; the positive and negative window sizes are +256 and -256 respectively, relative to the current maximum sequence number received. No

configuration is required to activate this feature; the window sizes are hard-coded and nonconfigurable.

## Intercepting IPv6 Flows

Starting with Junos OS Release 11.4, Junos Capture Vision also supports intercepting IPv6 flows in M320, T320, T640, and T1600 routers with a Multiservices 400 or Multiservices 500 PIC. Junos Capture Vision can intercept passively monitored IPv6 traffic only. All support for IPv4 interception remains the same. The interception of IPv6 traffic happens in the same way the filters capture IPv4 flows. With the introduction of IPv6 interception, both IPv4 and IPv6 filters can coexist. The mediation device, however, cannot be located in an IPv6 network.

Junos Capture Vision does not support interception of VPLS and MPLS traffic. The application cannot intercept Address Resolution Protocol (ARP) or other Layer 2 exception packets. The interception filter can be configured to timeout based on factors like total time (seconds), idle time (seconds), total packets or total data transmitted (bytes).

- Related Documentation**
- [Configuring Junos Capture Vision on page 849](#)
  - [Example: Configuring Junos Capture Vision on page 855](#)

---

## Configuring Junos Capture Vision

This section describes the following tasks for configuring Junos Capture Vision:

- [Configuring the Capture Group on page 849](#)
- [Configuring the Content Destination on page 850](#)
- [Configuring the Control Source on page 851](#)
- [Configuring the DFC PIC Interface on page 852](#)
- [Configuring the Firewall Filter on page 853](#)
- [Configuring System Logging on page 853](#)
- [Configuring Tracing Options for Junos Capture Vision Events on page 854](#)
- [Configuring Thresholds on page 854](#)
- [Limiting the Number of Duplicates of a Packet on page 855](#)

### Configuring the Capture Group

A capture group defines a profile of Junos Capture Vision configuration information. The static configuration includes information about control sources, content destinations, and notification destinations. Dynamic configuration is added through interaction with control sources using a control protocol.

To configure a capture group, include the **capture-group** statement at the **[edit services dynamic-flow-capture]** hierarchy level:

```
capture-group client-name {
 content-destination identifier {
 address address;
 }
}
```

```

 hard-limit bandwidth;
 hard-limit-target bandwidth;
 soft-limit bandwidth;
 soft-limit-clear bandwidth;
 ttl hops;
}
control-source identifier {
 allowed-destinations [destinations];
 minimum-priority value;
 no-syslog;
 notification-targets address port port-number;
 service-port port-number;
 shared-key value;
 source-addresses [addresses];
}
duplicates-dropped-periodicity seconds;
input-packet-rate-threshold rate;
interfaces interface-name;
max-duplicates number;
pic-memory-threshold percentage percentage;
}

```

To specify the **capture-group**, assign it a unique **client-name** that associates the information with the requesting control sources.

## Configuring the Content Destination

You must specify a destination for the packets that match DFC PIC filter criteria. To configure the content destination, include the **content-destination** statement at the [edit **services dynamic-flow-capture capture-group client-name**] hierarchy level:

```

content-destination identifier {
 address address;
 hard-limit bandwidth;
 hard-limit-target bandwidth;
 soft-limit bandwidth;
 soft-limit-clear bandwidth;
 ttl hops;
}

```

Assign the **content-destination** a unique **identifier**. You must also specify its IP address and you can optionally include additional settings:

- **address**—The DFC PIC interface appends an IP header with this destination address on the matched packet (with its own IP header and contents intact) and sends it out to the content destination.
- **ttl**—The time-to-live (TTL) value for the IP-IP header. By default, the TTL value is 255. Its range is 0 through 255.
- **Congestion thresholds**—You can specify per-content destination bandwidth limits that control the amount of traffic produced by the DFC PIC during periods of congestion. The thresholds are arranged in two pairs: **hard-limit** and **hard-limit-target**, and **soft-limit** and **soft-limit-clear**. You can optionally include one or both of these paired settings. All four settings are 10-second average bandwidth values in bits per second. Typically

**soft-limit-clear** < **soft-limit** < **hard-limit-target** < **hard-limit**. When the content bandwidth exceeds the **soft-limit** setting:

1. A congestion notification message is sent to each control source of the criteria that point to this content destination
2. If the control source is configured for **syslog**, a system log message is generated.
3. A latch is set, indicating that the control sources have been notified. No additional notification messages are sent until the latch is cleared, when the bandwidth falls below the **soft-limit-clear** value.

When the bandwidth exceeds the **hard-limit** value:

1. Junos Capture Vision begins deleting criteria until the bandwidth falls below the **hard-limit-target** value.
2. For each criterion deleted, a CongestionDelete notification is sent to the control source for that criterion.
3. If the control source is configured for **syslog**, a log message is generated.

The application evaluates criteria for deletion using the following data:

- **Priority**—Lower priority criteria are purged first, after adjusting for control source minimum priority.
- **Bandwidth**—Higher bandwidth criteria are purged first.
- **Timestamp**—The more recent criteria are purged first.

## Configuring the Control Source

You configure information about the control source, including allowed source addresses and destinations and authentication key values. To configure the control source information, include the **control-source** statement at the **[edit services dynamic-flow-capture capture-group *client-name*]** hierarchy level:

```
control-source identifier {
 allowed-destinations [destination-identifiers];
 minimum-priority value;
 no-syslog;
 notification-targets address port port-number;
 service-port port-number;
 shared-key value;
 source-addresses [addresses];
}
```

Assign the **control-source** statement a unique ***identifier***. You can also include values for the following statements:

- **allowed-destinations**—One or more content destination identifiers to which this control source can request that matched data be sent in its control protocol requests. If you do not specify any content destinations, all available destinations are allowed.
- **minimum-priority**—Value assigned to the control source that is added to the priority of the criteria in the DTCP ADD request to determine the total priority for the criteria. The

lower the value, the higher the priority. By default, **minimum-priority** has a value of 0 and the allowed range is 0 through 254.

- **notification-targets**—One or more destinations to which the DFC PIC interface can log information about control protocol-related events and other events such as PIC bootup messages. You configure each **notification-target** entry with an IP **address** value and a User Datagram Protocol (UDP) **port** number.
- **service-port**—UDP port number to which the control protocol requests are directed. Control protocol requests that are not directed to this port are discarded by DFC PIC interfaces.
- **shared-key**—20-byte authentication key value shared between the control source and the DFC PIC monitoring platform.
- **source-addresses**—One or more allowed IP addresses from which the control source can send control protocol requests to the DFC PIC monitoring platform. These are /32 addresses.

## Configuring the DFC PIC Interface

You specify the interface that interacts with the control sources configured in the same capture group. A Monitoring Services III PIC can belong to only one capture group, and you can configure only one PIC for each group.

To configure a DFC PIC interface, include the **interfaces** statement at the **[edit services dynamic-flow-capture capture-group client-name]** hierarchy level:

```
interfaces interface-name;
```

You specify DFC interfaces using the **dfc-** identifier at the **[edit interfaces]** hierarchy level. You must specify three logical units on each DFC PIC interface, numbered 0, 1, and 2. You cannot configure any other logical interfaces.

- **unit 0** processes control protocol requests and responses.
- **unit 1** receives monitored data.
- **unit 2** transmits the matched packets to the destination address.

The following example shows the configuration necessary to set up a DFC PIC interface and intercept both IPv4 and IPv6 traffic:

```
[edit interfaces dfc-0/0/0]
unit 0 {
 family inet {
 filter {
 output high; #Firewall filter to route control packets
 # through 'network-control' forwarding class. Control packets
 # are loss sensitive.
 }
 }
 address 10.1.0.0/32 { # DFC PIC address
 destination 10.36.100.1; # DFC PIC address used by
 # the control source to correspond with the
 # monitoring platform
 }
}
```

```

 }
 }
 unit 1 { # receive data packets on this logical interface
 family inet; # receive IPv4 traffic for interception
 family inet6; # receive IPv6 traffic for interception
 }
 unit 2 { # send out copies of matched packets on this logical interface
 family inet;
 }
}

```

In addition, you must configure Junos Capture Vision to run on the DFC PIC in the correct chassis location. The following example shows this configuration at the **[edit chassis]** hierarchy level:

```

fpc 0 {
 pic 0 {
 monitoring-services application dynamic-flow-capture;
 }
}

```

For more information on configuring chassis properties, see the *Junos OS Administration Library for Routing Devices*.

## Configuring the Firewall Filter

You can specify the firewall filter to route control packets through the network control forwarding class. The control packets are loss sensitive. To configure the firewall filter, include the following statements at the **[edit]** hierarchy level:

```

firewall {
 family inet {
 filter high {
 term all {
 then forwarding-class network-control;
 }
 }
 }
}

```

## Configuring System Logging

By default, control protocol activity is logged as a separate system log facility, **dfc**. To modify the filename or level at which control protocol activity is recorded, include the following statements at the **[edit syslog]** hierarchy level:

```

file dfc.log {
 dfc any;
}

```

To cancel logging, include the **no-syslog** statement at the **[edit services dynamic-flow-capture capture-group *client-name* control-source *identifier*]** hierarchy level:

```

no-syslog;

```



**NOTE:** Junos Capture Vision (dfc-) interface supports up to 10,000 filter criteria. When more than 10,000 filters are added to the interface, the filters are accepted, but system log messages are generated indicating that the filter is full.

---

## Configuring Tracing Options for Junos Capture Vision Events

You can enable tracing options for Junos Capture Vision events by including the **traceoptions** statement at the **[edit services dynamic-flow-capture]** hierarchy level.

When you include the **traceoptions** configuration, you can also specify the trace file name, maximum number of trace files, the maximum size of trace files, and whether the trace file can be read by all users or not.

To enable tracing options for Junos Capture Vision events, include the following configuration at the **[edit services dynamic-flow-capture]** hierarchy level:

```
traceoptions{
 file filename <files number> <size size> <world-readable | non-world-readable>;
}
```

To disable tracing for Junos Capture Vision events, delete the **traceoptions** configuration from the **[edit services dynamic-flow-capture]** hierarchy level.



**NOTE:** In Junos OS releases earlier than 9.2R1, tracing of Junos Capture Vision was enabled by default, and the logs were saved to the `/var/log/dfcd` directory.

---

## Configuring Thresholds

You can optionally specify threshold values for the following situations in which warning messages will be recorded in the system log:

- Input packet rate to the DFC PIC interfaces
- Memory usage on the DFC PIC interfaces

To configure threshold values, include the **input-packet-rate-threshold** or **pic-memory-threshold** statements at the **[edit services dynamic-flow-capture capture-group *client-name*]** hierarchy level:

```
input-packet-rate-threshold rate;
pic-memory-threshold percentage percentage;
```

If these statements are not configured, no threshold messages are logged. The threshold settings are configured for the capture group as a whole.

The range of configurable values for the **input-packet-rate-threshold** statement is 0 through 1 Mpps. The PIC calibrates the value accordingly; the Monitoring Services III PIC caps the threshold value at 300 Kpps and the Multiservices 400 PIC uses the full



configured value. The range of values for the **pic-memory-threshold** statement is 0 to 100 percent.

## Limiting the Number of Duplicates of a Packet

You can optionally specify the maximum number of duplicate packets the DFC PIC is allowed to generate from a single input packet. This limitation is intended to reduce the load on the PIC when packets are sent to multiple destinations. When the maximum number is reached, the duplicates are sent to the destinations with the highest criteria class priority. Within classes of equal priority, criteria having earlier timestamps are selected first.

To configure this limitation, include the **max-duplicates** statement at the **[edit services dynamic-flow-capture capture-group *client-name*]** hierarchy level:

```
max-duplicates number;
```

You can also apply the limitation on a global basis for the DFC PIC by including the **g-max-duplicates** statement at the **[edit services dynamic-flow-capture]** hierarchy level:

```
g-max-duplicates number;
```

By default, the maximum number of duplicates is set to 3. The range of allowed values is 1 through 64. A setting for **max-duplicates** for an individual capture-group overrides the global setting.

In addition, you can specify the frequency with which the application sends notifications to the affected control sources that duplicates are being dropped because the threshold has been reached. You configure this setting at the same levels as the maximum duplicates settings, by including the **duplicates-dropped-periodicity** statement at the **[edit services dynamic-flow-capture capture-group *client-name*]** hierarchy level or the **g-duplicates-dropped-periodicity** statement at the **[edit services dynamic-flow-capture]** hierarchy level:

```
duplicates-dropped-periodicity seconds;
g-duplicates-dropped-periodicity seconds;
```

As with the **g-max-duplicates** statement, the **g-duplicates-dropped-periodicity** statement applies the setting globally for the application and is overridden by a setting applied at the capture-group level. By default, the frequency for sending notifications is 30 seconds.

- Related Documentation**
- [Understanding Junos Capture Vision on page 847](#)
  - [Example: Configuring Junos Capture Vision on page 855](#)

---

## Example: Configuring Junos Capture Vision

The following example includes all parts of a complete Junos Capture Vision configuration.

Configure the Junos Capture Vision PIC interface:

```
[edit interfaces dfc-0/0/0]
unit 0 {
 family inet {
```

```
filter {
 output high; #Firewall filter to route control packets
 # through 'network-control' forwarding class. Control packets
 # are loss sensitive.
}
address 10.1.0.0/32 { # DFC PIC address
 destination 10.36.100.1; # DFC PIC address used by
 # the control source to correspond with the
 # monitoring platform
}
}
unit 1 { # receive data packets on this logical interface
 family inet;
 family inet6;
}
unit 2 { # send out copies of matched packets on this logical interface
 family inet;
}
```

Configure the capture group:

```
services dynamic-flow-capture {
 capture-group g1 {
 interfaces dfc-0/0/0;
 input-packet-rate-threshold 90k;
 pic-memory-threshold percentage 80;
 control-source cs1 {
 source-addresses 10.36.41.1;
 service-port 2400;
 notification-targets {
 10.36.41.1 port 2100;
 }
 shared-key "9ASxdsYoX7wg4aHk";
 allowed-destinations cd1;
 }
 content-destination cd1 {
 address 10.36.70.2;
 ttl 244;
 }
 }
}
```

Configure filter-based forwarding (FBF) to the Junos Capture Vision PIC interface, logical unit 1.

For more information about configuring passive monitoring interfaces, see [“Enabling Passive Flow Monitoring” on page 830](#).

```
interfaces so-1/2/0 {
 encapsulation ppp;
 unit 0 {
 passive-monitor-mode;
 family inet {
 filter {
 input catch;
 }
 }
 }
}
```

```

 }
 }
}

```

Configure the firewall filter:

```

firewall {
 filter catch {
 interface-specific;
 term def {
 then {
 count counter;
 routing-instance fbf_inst;
 }
 }
 }
 family inet {
 filter high {
 term all {
 then forwarding-class network-control;
 }
 }
 }
}

```

Configure a forwarding routing instance. The next hop points specifically to the logical interface corresponding to **unit 1**, because only this particular logical unit is expected to relay monitored data to the Junos Capture Vision PIC.

```

routing-instances fbf_inst {
 instance-type forwarding;
 routing-options {
 static {
 route 0.0.0.0/0 next-hop dfc-0/0/0.1;
 }
 }
}

```

Configure routing table groups:

```

[edit]
routing-options {
 interface-routes {
 rib-group inet common;
 }
 rib-groups {
 common {
 import-rib [inet.0 fbf_inst.inet.0];
 }
 }
 forwarding-table {
 export pplb;
 }
}

```

Configure interfaces to the control source and content destination:

```

interfaces fe-4/1/2 {
 description "to cs1 from dfc";
}

```

```
unit 0 {
 family inet {
 address 10.36.41.2/30;
 }
}
}
interfaces ge-7/0/0 {
 description "to cd1 from dfc";
 unit 0 {
 family inet {
 address 10.36.70.1/30;
 }
 }
}
```

- Related Documentation**
- [Understanding Junos Capture Vision on page 847](#)
  - [Configuring Junos Capture Vision on page 849](#)

# Detecting Threats and Intercepting Flows Using Junos Packet Vision

- [Understanding Junos Packet Vision on page 859](#)
- [Junos Packet Vision Architecture on page 860](#)
- [Configuring Junos Packet Vision on page 861](#)
- [Configuring FlowTapLite on page 864](#)
- [Examples: Configuring Junos Packet Vision on page 865](#)

## Understanding Junos Packet Vision

---

Junos Capture Vision (previously known as dynamic flow capture) enables you to capture packet flows on the basis of dynamic filtering criteria, using Dynamic Tasking Control Protocol (DTCP) requests. Junos Packet Vision is a Junos OS application that performs lawful intercept of packet flows, using Dynamic Tasking Control Protocol (DTCP). The application extends the use of DTCP to intercept IPv4 and IPv6 packets in an active monitoring router and send a copy of packets that match filter criteria to one or more content destinations. Junos Packet Vision was previously known as flow-tap application.

Junos Packet Vision data can be used in the following applications:

- Flexible trend analysis for detection of new security threats
- Lawful intercept

Junos Packet Vision is supported on M Series and T Series routers, except M160 and TX Matrix routers. Junos Packet Vision filters are applied on all IPv4 traffic and do not add any perceptible delay in the forwarding path. Junos Packet Vision filters can also be applied on IPv6 traffic. For security, filters installed by one client are not visible to others and the CLI configuration does not reveal the identity of the monitored target. A lighter version of the application is supported on MX Series routers only.

### Related Documentation

- [Junos Packet Vision Architecture on page 860](#)
- [Configuring Junos Packet Vision on page 861](#)
- [Configuring FlowTapLite on page 864](#)
- [Examples: Configuring Junos Packet Vision on page 865](#)

## Junos Packet Vision Architecture

---

The Junos Packet Vision (previously known as Flow-Tap) architecture consists of one or more *mediation devices* that send requests to a Juniper Networks router to monitor incoming data and forward any packets that match specific filter criteria to a set of one or more *content destinations*:

- **Mediation device**—A client that monitors electronic data or voice transfer over the network. The mediation device sends filter requests to the Juniper Networks router using the DTCP. The clients are not identified for security reasons, but have permissions defined by a set of special login classes. Each system can support up to 16 different mediation devices for each user, up to a maximum of 64 mediation devices for the whole system.
- **Monitoring platform**—An M Series or T Series router containing one or more Adaptive Services (AS) or Multiservices PICs, which are configured to support the Junos Packet Vision application. The monitoring platform processes the requests from the mediation devices, applies the dynamic filters, monitors incoming data flows, and sends the matched packets to the appropriate content destinations.
- **Content destination**—Recipient of the matched packets from the monitoring platform. Typically the matched packets are sent using an IP Security (IPsec) tunnel from the monitoring platform to another router connected to the content destination. The content destination and the mediation device can be physically located on the same host. For more information about IPsec tunnels, see *Junos VPN Site Secure*.
- **Dynamic filters**—Firewall filters automatically generated by the Packet Forwarding Engine and applied to all routing instances. Each term in the filter includes a **flow-tap** action that is similar to the existing **sample** or **port-mirroring** actions. As long as one of the filter terms matches an incoming packet, the router copies the packet and forwards it to the Adaptive Services or Multiservices PIC that is configured for Junos Packet Vision service. The Adaptive Services or Multiservices PIC runs the packet through the client filters and sends a copy to each matching content destination.

Following is a sample filter configuration; note that it is dynamically generated by the router (no user configuration is required):

```
filter combined_LEA_filter {
 term LEA1_filter {
 from {
 source-address 1.2.3.4;
 destination-address 3.4.5.6;
 }
 then {
 flow-tap;
 }
 }
 term LEA2_filter {
 from {
 source-address 10.1.1.1;
 source-port 23;
 }
 }
}
```

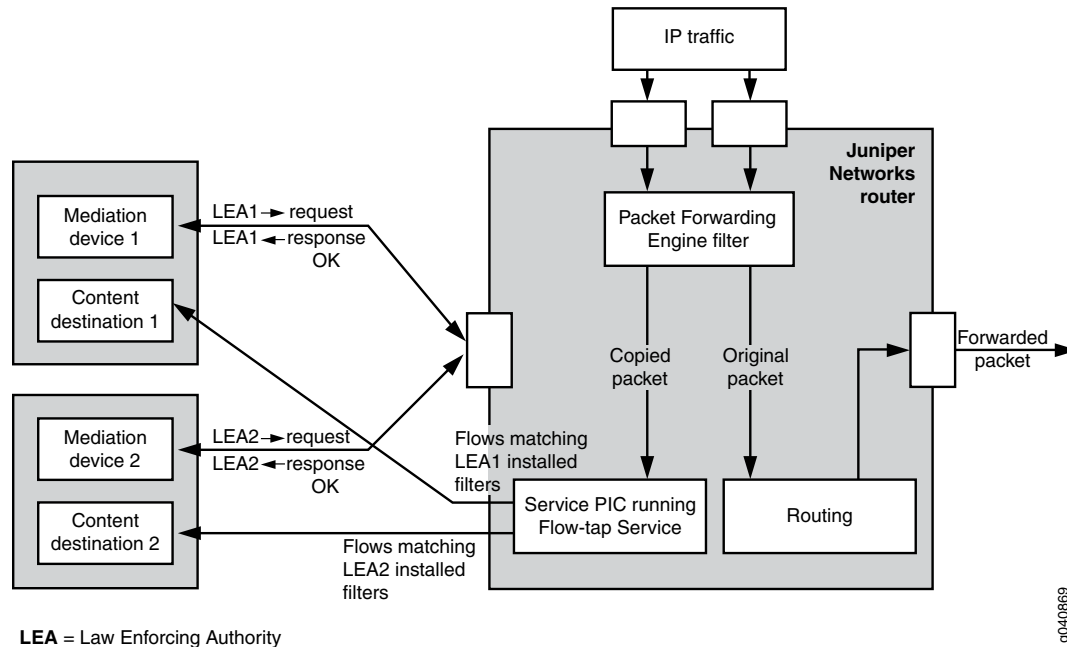
```

 then {
 flow-tap;
 }
 }
}

```

Figure 33 on page 861 shows a sample topology that uses two mediation devices and two content destinations.

Figure 33: Junos Packet Vision Topology



#### Related Documentation

- [Understanding Junos Packet Vision on page 859](#)
- [\[edit services flow-tap\] Hierarchy Level on page 1631](#)
- [Configuring Junos Packet Vision on page 861](#)
- [Examples: Configuring Junos Packet Vision on page 865](#)

## Configuring Junos Packet Vision

This topic explains Junos Packet Vision (previously known as Flow-Tap) configuration, and contains the following sections:

- [Configuring the Junos Packet Vision Interface on page 861](#)
- [Strengthening Junos Packet Vision Security on page 862](#)
- [Restrictions on Junos Packet Vision Services on page 863](#)

### Configuring the Junos Packet Vision Interface

To configure an adaptive services interface for flow-tap service, include the **interface** statement at the **[edit services flow-tap]** hierarchy level:

```
interface sp-fpc/pic/port.unit-number;
```

You can assign any Adaptive Services or Multiservices PIC in the active monitoring router for Junos Packet Vision, and use any logical unit on the PIC.

You can specify the type of traffic for which you want to apply the Junos Packet Vision service by including the **family inet | inet6** statement. If the **family** statement is not included, the Junos Packet Vision service is, by default, applied to the IPv4 traffic. To apply Junos Packet Vision service to IPv6 traffic, you must include the **family inet6** statement in the configuration. To enable the Junos Packet Vision service for IPv4 and IPv6 traffic, you must explicitly configure the **family** statement for both **inet** and **inet6** families.



**NOTE:** You cannot configure Junos Capture Vision (previously known as dynamic flow capture) and Junos Packet Vision services on the same router simultaneously.

You must also configure the logical interface at the **[edit interfaces]** hierarchy level:

```
interface sp-fpc/pic/port {
 unit logical-unit-number {
 family inet;
 family inet6;
 }
}
```



**NOTE:** If you do not include the **family inet6** statement in the configuration, IPv6 flows will not be intercepted.

## Strengthening Junos Packet Vision Security

You can add an extra level of security to Dynamic Tasking Control Protocol (DTCP) transactions between the mediation device and the router by enabling DTCP sessions on top of the SSH layer. To configure SSH settings, include the **flow-tap-dtcp** statement at the **[edit system services]** hierarchy level:

```
flow-tap-dtcp {
 ssh {
 connection-limit value;
 rate-limit value;
 }
}
```

To configure client permissions for viewing and modifying Junos Packet Vision configurations and for receiving tapped traffic, include the **permissions** statement at the **[edit system login class class-name]** hierarchy level:

```
permissions [permissions];
```

The permissions needed to use Junos Packet Vision features are as follows:

- **flow-tap**—Can view Junos Packet Vision configuration



- **flow-tap-control**—Can modify Junos Packet Vision configuration
- **flow-tap-operation**—Can tap flows

You can also specify user permissions on a RADIUS server, for example:

```
Bob Auth-Type := Local, User-Password = = "abc123"
Juniper-User-Permissions = "flow-tap-operation"
```

For details on **[edit system]** and RADIUS configuration, see the *Junos OS Administration Library for Routing Devices*.

## Restrictions on Junos Packet Vision Services

The following restrictions apply to Junos Packet Vision services:

- You cannot configure Junos Capture Vision and Junos Packet Vision features on the same router simultaneously.
- On routers that support LMNR-based FPCs, you cannot configure the Junos Packet Vision for IPv6 along with port mirroring or sampling of IPv6 traffic. This restriction applies even if the router does not have any LMNR-based FPC installed in it. However, there is no restriction on configuring Junos Packet Vision on routers that are configured for port mirroring or sampling of IPv4 traffic.
- Junos Packet Vision does not support interception of MPLS and virtual private LAN service (VPLS).
- Junos Packet Vision cannot intercept Address Resolution Protocol (ARP) and other Layer 2 exceptions.
- IPv4 and IPv6 intercept filters can coexist on a system, subject to a combined maximum of 100 filters.
- When Junos Capture Vision process or the Adaptive Services or Multiservices PIC configured for Junos Packet Vision restarts, all filters are deleted and the mediation devices are disconnected.
- Only the first fragment of an IPv4 fragmented packet stream is sent to the content destination.
- Port mirroring might not work in conjunction with Junos Packet Vision.
- Running the Junos Packet Vision over an IPsec tunnel on the same router can cause packet loops and is not supported.
- M10i routers do not support the standard Junos Packet Vision, but do support FlowTapLite (see [“Configuring FlowTapLite” on page 864](#)). Junos Packet Vision and FlowTapLite cannot be configured simultaneously on the same chassis.
- PIC-based flow-tap is not supported on M7i and M10i routers equipped with an Enhanced Compact Forwarding Engine Board (CFEB-E).
- You cannot configure Junos Packet Vision on channelized interfaces.

### Related Documentation

- [Configuring FlowTapLite on page 864](#)

## Configuring FlowTapLite

A lighter version of the flow-tap application is available on MX Series routers and also on M320 routers with Enhanced III Flexible PIC Concentrators (FPCs). All of the functionality resides in the Packet Forwarding Engine rather than a service PIC or Dense Port Concentrator (DPC).



**NOTE:** On M320 routers only, if the replacement of FPCs results in a mode change, you must restart the dynamic flow capture process manually by disabling and then re-enabling the CLI configuration.

FlowTapLite uses the same DTCP-SSH architecture to install the Dynamic Tasking Control Protocol (DTCP) filters and authenticate the users as the original flow-tap application and supports up to 3000 filters per chassis.



**NOTE:** The original flow-tap application and FlowTapLite cannot be used at the same time.

To configure FlowTapLite, include the **flow-tap** statement at the **[edit services]** hierarchy level:

```
flow-tap {
 tunnel-interface interface-name;
}
```

For the Packet Forwarding Engine to encapsulate the intercepted packet, it must send the packet to a tunnel logical (**vt-**) interface. You need to allocate a tunnel interface and assign it to the dynamic flow capture process for FlowTapLite to use. To create the tunnel interface, include the following configuration:

```
chassis {
 fpc number {
 pic number {
 tunnel-services {
 bandwidth (1g | 10g);
 }
 }
 }
}
```



**NOTE:** Currently FlowTapLite supports only one tunnel interface per instance.

For more information about this configuration, see the *Junos OS Administration Library for Routing Devices*.

To configure the logical interfaces and assign them to the dynamic flow capture process, include the following configuration:

```

interfaces {
 vt-fpc/pic/port {
 unit 0 {
 family inet;
 family inet6;
 }
 }
}

```



**NOTE:** If a service PIC or DPC is available, you can use its tunnel interface for the same purpose.



**NOTE:** If you do not include the `family inet6` statement in the configuration, IPv6 flows will not be intercepted.



**NOTE:** With FlowTapLite configured and traceoptions enabled, if you add more than two content destinations by including the X-JTAP-CDEST-DEST-ADDRESS line in the Dynamic Tasking Control Protocol (DTCP) parameter file and initiate a DTCP session by sending a DTCP ADD message, a '400 BAD request' message is received. Although you can specify more than two content destinations in the DTCP file that is sent from the mediation device, this error message occurs when the DTCP ADD message is sent. This behavior is expected with more than two content destinations. You must specify only two content destinations per DTCP ADD message.

#### Related Documentation

- [Understanding Junos Packet Vision on page 859](#)
- [\[edit services flow-tap\] Hierarchy Level on page 1631](#)
- [Configuring Junos Packet Vision on page 861](#)
- [Examples: Configuring Junos Packet Vision on page 865](#)

## Examples: Configuring Junos Packet Vision

The following example shows all parts of a complete Junos Packet Vision configuration with IPv4 and IPv6 flow intercepts



**NOTE:** The following example applies only to M Series and T Series routers, except M160 and TX Matrix routers. For MX Series routers, because the flow-tap application resides in the Packet Forwarding Engine rather than a service PIC or Dense Port Concentrator (DPC), the Packet Forwarding Engine must send the packet to a tunnel logical (vt-) interface to encapsulate the intercepted packet. In such a scenario, you need to allocate a tunnel interface and assign it to the dynamic flow capture process for FlowTapLite to use.

```
services {
 flow-tap {
 interface sp-1/2/0.100;
 }
}
interfaces {
 sp-1/2/0 {
 unit 100 {
 family inet;
 family inet6;
 }
 }
}
system {
 services {
 flow-tap-dtcp {
 ssh {
 connection-limit 5;
 rate-limit 5;
 }
 }
 }
 login {
 class ft-class {
 permissions flow-tap-operation;
 }
 user ft-user1 {
 class ft-class;
 authentication {
 encrypted-password "xxxx";
 }
 }
 }
}
```

The following example shows a FlowTapLite configuration that intercepts IPv4 and IPv6 flows:

```
system {
 login {
 class flowtap {
 permissions flow-tap-operation;
 }
 user ftap {
 uid 2000;
 class flowtap;
 authentication {
 encrypted-password "1nZfwNn4L$TWi/oxFwFZyOyyxN/87Jv0"; ##
 SECRET-DATA
 }
 }
 }
 services {
 flow-tap-dtcp {
 ssh;
 }
 }
}
```

```
 }
 }
 chassis {
 fpc 0 {
 pic 0 {
 tunnel-services {
 bandwidth 10g;
 }
 }
 }
 }
 interfaces {
 vt-0/0/0 {
 unit 0 {
 family inet;
 family inet6;
 }
 }
 }
 services {
 flow-tap {
 tunnel-interface vt-0/0/0.0;
 }
 }
}
```

**Related  
Documentation**

- [Understanding Junos Packet Vision on page 859](#)
- [\[edit services flow-tap\] Hierarchy Level on page 1631](#)
- [Configuring Junos Packet Vision on page 861](#)
- [Configuring FlowTapLite on page 864](#)



## PART 16

# Sampling, Discard Accounting, and Port Mirroring Services

- [Sampling Data Using Traffic Sampling and Discard Accounting on page 871](#)
- [Sampling Data Using Inline Sampling on page 885](#)
- [Sampling Data Using Flow Aggregation on page 897](#)
- [Sending Packets for Analysis Using Port Mirroring on page 931](#)





# Sampling Data Using Traffic Sampling and Discard Accounting

- [Configuring Traffic Sampling on page 871](#)
- [Sampling Instance Configuration on page 881](#)
- [Configuring Discard Accounting on page 883](#)

## Configuring Traffic Sampling

---

Traffic sampling enables you to copy traffic to a Physical Interface Card (PIC) that performs flow accounting while the router forwards the packet to its original destination. You can configure the router to perform sampling in either of two locations:

- On the Routing Engine, using the sampled process. To select this method, use a filter (input or output) with a matching term that contains the **then sample** statement.
- On the Monitoring Services, Adaptive Services, or Multiservices PIC.



**NOTE:** Routing Engine based sampling is not supported on VPN routing and forwarding (VRF) instances.

The following sections provide configuration instructions for traffic sampling:

- [Configuring Firewall Filter for Traffic Sampling on page 871](#)
- [Configuring Traffic Sampling on a Logical Interface on page 873](#)
- [Disabling Traffic Sampling on page 874](#)
- [Sampling Once on page 874](#)
- [Preserving Prerewrite ToS Value for Egress Sampled or Mirrored Packets on page 875](#)
- [Configuring Traffic Sampling Output on page 876](#)
- [Tracing Traffic Sampling Operations on page 878](#)
- [Traffic Sampling Examples on page 878](#)

## Configuring Firewall Filter for Traffic Sampling

To configure firewall filter for traffic sampling, you must perform the following tasks:

- Create a firewall filter to apply to the logical interfaces being sampled by including the **filter** statement at the **[edit firewall family *family-name*]** hierarchy level. In the filter **then** statement, you must specify the action modifier **sample** and the action **accept**.

```
filter filter-name {
 term term-name {
 then {
 sample;
 accept;
 }
 }
}
```

For more information about firewall filter actions and action modifiers, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

- Apply the filter to the interfaces on which you want to sample traffic by including the **address** and **filter** statements at the **[edit interfaces *interface-name* unit *logical-unit-number* family *family-name*]** hierarchy level:

```
address address {
}
filter {
 input filter-name;
}
```

The following prerequisites apply to M, MX, and T Series routers when you configure traffic sampling on interfaces and in firewall filters:

- If you configure a sample action in a firewall filter for an inet or inet6 family on an interface without configuring the forwarding-options settings, operational problems might occur if you also configure port mirroring or flow-tap functionalities. In such a scenario, all the packets that match the firewall filter are incorrectly sent to the service PIC.
- If you include the **then sample** statement at the **[edit firewall family inet filter *filter-name* term *term-name*]** hierarchy level to specify a sample action in a firewall filter for IPv4 packets, you must also include the **family inet** statement at the **[edit forwarding-options sampling]** hierarchy level or the **instance *instance-name* family inet** statement at the **[edit forwarding-options sampling]** hierarchy level. Similarly, if you include the **then sample** statement at the **[edit firewall family inet6 filter *filter-name* term *term-name*]** hierarchy level to specify a sample action in a firewall filter for IPv6 packets, you must also include **family inet6** statement at the **[edit forwarding-options sampling]** hierarchy level or the **instance *instance-name* family inet6** statement at the **[edit forwarding-options sampling]** hierarchy level. Otherwise, a commit error occurs when you attempt to commit the configuration.
- Also, if you configure traffic sampling on a logical interface by including the sampling input or sampling output statements at the **[edit interface *interface-name* unit *logical-unit-number*]** hierarchy level, you must also include the **family inet | inet6** statement at the **[edit forwarding-options sampling]** hierarchy level, or the **instance *instance-name* family inet | inet6** statement at the **[edit forwarding-options sampling]** hierarchy level.

## Configuring Traffic Sampling on a Logical Interface

To configure traffic sampling on any logical interface, enable sampling and specify a non zero sampling rate by including the sampling statement at the **[edit forwarding-options]** hierarchy level:

```
sampling {
 input {
 rate number;
 run-length number;
 max-packets-per-second number;
 maximum-packet-length bytes;
 }
}
```

When you use Routing Engine-based sampling, specify the threshold traffic value by including the **max-packets-per-second** statement. The value is the maximum number of packets to be sampled, beyond which the sampling mechanism begins dropping packets. The range is from 0 through 65,535. A value of 0 instructs the Packet Forwarding Engine not to sample any packets. The default value is 1000.



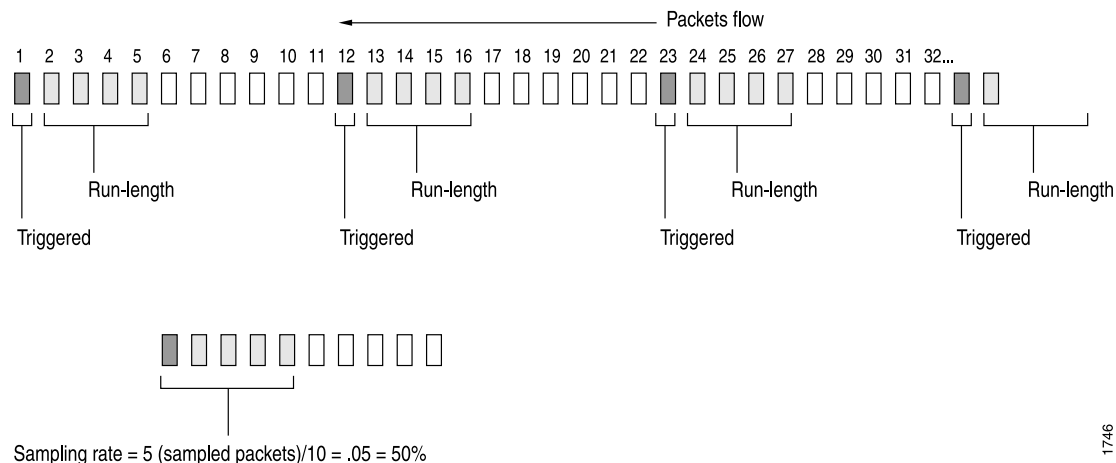
**NOTE:** When you configure active monitoring and specify a Monitoring Services, Adaptive Services, or Multiservices PIC in the output statement, the **max-packets-per-second** value is ignored.

Specify the sampling rate by setting the values for **rate** and **run-length** (see [Figure 34 on page 873](#)).

**Figure 34: Configuring Sampling Rate**

### Rate and Run-length

Case #1 Rate = 10, run-length = 4



1746



**NOTE:** If PIC-based flow monitoring is enabled on an *ms-fpc/pic/port.logical-unit* interface, a commit check error occurs when you attempt to configure ingress traffic sampling on that interface. This error occurs because a combination of ingress sampling and PIC-based flow monitoring operations on an ms- logical interface causes undesired flow monitoring behavior and might result in repeated sampling of a single packet. You must not configure ingress sampling on ms- logical interfaces on which PIC-based flow monitoring is enabled.

The **rate** statement specifies the ratio of packets to be sampled. For example, if you configure a rate of 10, x number of packets out of every 10 is sampled, where  $x = \text{run length} + 1$ . By default, the rate is 0, which means that no traffic is sampled.

The **run-length** statement specifies the number of matching packets to sample following the initial one-packet trigger event. By default, the run length is 0, which means that no more traffic is sampled after the trigger event. The range is from 0 through 20. Configuring a run length greater than 0 allows you to sample packets following those already being sampled.



**NOTE:** The **run-length** and **maximum-packet-length** configuration statements are not supported on MX80 routers.

If you do not include the **input** statement, sampling is disabled.

To collect the sampled packets in a file, include the **file** statement at the **[edit forwarding-options sampling output]** hierarchy level. Output file formats are discussed later in the chapter.

## Disabling Traffic Sampling

To explicitly disable traffic sampling on the router, include the **disable** statement at the **[edit forwarding-options sampling]** hierarchy level:

```
disable;
```

## Sampling Once

To explicitly sample a packet for active monitoring only once, include the **sample-once** statement at the **[edit forwarding-options sampling]** hierarchy level:

```
sample-once;
```

Setting this option avoids duplication of packets in cases where sampling is enabled at both the ingress and egress interfaces and simplifies analysis of the sampled traffic.

## Preserving Prerewrite ToS Value for Egress Sampled or Mirrored Packets

To preserve the prenormalized type-of-service (ToS) value in egress sampled or mirrored packets, include the **pre-rewrite-tos** statement at the **[edit forwarding-options sampling]** hierarchy level.

On MPC-based interfaces, you can configure ToS rewrite either using class-of-service (CoS) configuration by including the **rewrite-rules dscp rule\_name** statement at the **[edit class-of-service interfaces interface-name unit logical-unit-number]** hierarchy level or using firewall filter configuration by including the **dscp** statement at the **[edit firewall family family-name filter filter-name term term-name then]** hierarchy level. If ToS rewrite is configured, the egress mirrored or sampled copies contain the post-rewrite ToS values by default. With the **pre-rewrite-tos** configuration, you can retain the prerewrite ToS value in the sampled or mirrored packets.



### NOTE:

- If ToS rewrite is configured on the egress interface by using both CoS and firewall filter configuration, and if the **pre-rewrite-tos** statement is also configured, then the egress sampled packets contain the DSCP value set using the firewall filter configuration. However, if the **pre-rewrite-tos** statement is not configured, the egress sampled packets contain the DSCP value set by the CoS configuration.
- With the **pre-rewrite-tos** statement, you can configure retaining prenormalization ToS values only for sampling done under family inet and family inet6.
- This feature cannot be configured at the **[edit logical-systems]** hierarchy level. It can be configured only at the global level under the forwarding-option configuration.
- When ToS rewrite is configured by using a firewall filter on both ingress and egress interfaces, the egress sampled packets contain the DSCP value set by the ingress ToS rewrite configuration if the **pre-rewrite-tos** statement is configured. However, if the **pre-rewrite-tos** statement is not configured, the egress sampled packets contain the DSCP value set by the ToS rewrite configuration for the egress firewall filter.
- If the **pre-rewrite-tos** statement is configured, and a deactivate or delete operation is performed at the **[edit forwarding-options]** hierarchy level, **pre-rewrite-tos** configuration still remains active. To disable the **pre-rewrite-tos** configuration for such a case, you must explicitly deactivate or delete the **pre-rewrite-tos** statement at the **[edit forwarding-options sampling]** hierarchy level before performing a deactivate or delete operation at the **[edit forwarding-options]** hierarchy level.

## Configuring Traffic Sampling Output

To configure traffic sampling output, include the following statements at the **[edit forwarding-options sampling family (inet | inet6 | mpls) output]** hierarchy level:

```
aggregate-export-interval seconds;
flow-active-timeout seconds;
flow-inactive-timeout seconds;
extension-service service-name;
flow-server hostname {
 aggregation {
 autonomous-system;
 destination-prefix;
 protocol-port;
 source-destination-prefix {
 caida-compliant;
 }
 source-prefix;
 }
 autonomous-system-type (origin | peer);
 (local-dump | no-local-dump);
 port port-number;
 source-address address;
 version format;
 version9 {
 template template-name;
 }
}
interface interface-name {
 engine-id number;
 engine-type number;
 source-address address;
}
file {
 disable;
 filename filename;
 files number;
 size bytes;
 (stamp | no-stamp);
 (world-readable | no-world-readable);
}
```

To configure inline flow monitoring on MX Series routers, include the **inline-jflow** statement at the **[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output]** hierarchy level. Inline sampling exclusively supports a new format called IP\_FIX that uses UDP as the transport protocol. When you configure inline sampling, you must include the **version-ipfix** statement at the **[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output flow-server address]** hierarchy level and also at the **[edit services flow-monitoring]** hierarchy level. For more information about configuring inline flow monitoring, see [“Configuring Inline Active flow Monitoring” on page 890](#).

To direct sampled traffic to a flow-monitoring interface, include the **interface** statement. The **engine-id** and **engine-type** statements specify the identity and type numbers of the

interface; they are dynamically generated based on the Flexible PIC Concentrator (FPC), PIC, and slot numbers and the chassis type. The **source-address** statement specifies the traffic source.

To configure flow sampling version 9 output, you need to include the **template** statement at the **[edit forwarding-options sampling output version9]** hierarchy level. For information on cflowd, see [“Enabling Flow Aggregation” on page 898](#).

The **aggregate-export-interval** statement is described in [“Configuring Discard Accounting” on page 883](#), and the **flow-active-timeout** and **flow-inactive-timeout** statements are described in [“Configuring Flow Monitoring” on page 818](#).

Traffic sampling results are automatically saved to a file in the **/var/tmp** directory. To collect the sampled packets in a file, include the **file** statement at the **[edit forwarding-options sampling family inet output]** hierarchy level:

```
file {
 disable;
 filename filename;
 files number;
 size bytes;
 (stamp | no-stamp);
 (world-readable | no-world-readable);
}
```

### Traffic Sampling Output Format

Traffic sampling output is saved to an ASCII text file. The following is an example of the traffic sampling output that is saved to a file in the **/var/tmp** directory. Each line in the output file contains information for one sampled packet. You can optionally display a timestamp for each line.

The column headers are repeated after each group of 1000 packets.

```
Apr 7 15:48:50
Time Dest Src Dest Src Proto TOS Pkt Intf IP TCP
 addr addr port port len num frag flags
Apr 7 15:48:54 192.168.9.194 192.168.9.195 0 0 1 0x0 84 8 0x0 0x0
Apr 7 15:48:55 192.168.9.194 192.168.9.195 0 0 1 0x0 84 8 0x0 0x0
Apr 7 15:48:56 192.168.9.194 192.168.9.195 0 0 1 0x0 84 8 0x0 0x0
Apr 7 15:48:57 192.168.9.194 192.168.9.195 0 0 1 0x0 84 8 0x0 0x0
Apr 7 15:48:58 192.168.9.194 192.168.9.195 0 0 1 0x0 84 8 0x0 0x0
```

To set the timestamp option for the file **my-sample**, enter the following:

```
[edit forwarding-options sampling output file]
user@host# set filename my-sample files 5 size 2m world-readable stamp;
```

Whenever you toggle the timestamp option, a new header is included in the file. If you set the **stamp** option, the **Time** field is displayed.

```
Apr 7 15:48:50
Time Dest Src Dest Src Proto TOS Pkt Intf IP TCP
addr addr port port len num frag flags
Feb 1 20:31:21
```

| # | Dest | Src  | Dest | Src  | Proto | TOS | Pkt | Intf | IP   | TCP   |
|---|------|------|------|------|-------|-----|-----|------|------|-------|
| # | addr | addr | port | port |       |     | len | num  | frag | flags |

## Tracing Traffic Sampling Operations

Tracing operations track all traffic sampling operations and record them in a log file in the `/var/log` directory. By default, this file is named `/var/log/sampled`. The default file size is 128K, and 10 files are created before the first one gets overwritten.

To trace traffic sampling operations, include the **traceoptions** statement at the **[edit forwarding-options sampling]** hierarchy level:

```
traceoptions {
 no-remote-trace;
 file filename <files number> <size bytes> <match expression> <world-readable |
 no-world-readable>;
}
```

## Traffic Sampling Examples

The following sections provide examples of configuring traffic sampling:

- [Example: Sampling a Single SONET/SDH Interface on page 878](#)
- [Example: Sampling All Traffic from a Single IP Address on page 879](#)
- [Example: Sampling All FTP Traffic on page 880](#)

---

### Example: Sampling a Single SONET/SDH Interface

The following configuration gathers statistical sampling information from a small percentage of all traffic on a single SONET/SDH interface and collects it in a file named **sonet-samples.txt**.

Create the filter:

```
[edit firewall family inet]
filter {
 input sample-sonet {
 then {
 sample;
 accept;
 }
 }
}
```

Apply the filter to the SONET/SDH interface:

```
[edit interfaces]
so-0/0/1 {
 unit 0 {
 family inet {
 filter {
 input sample-sonet;
 }
 address 10.127.68.254/32 {
 destination 172.16.74.7;
```



```

 }
 }
}

```

Finally, configure traffic sampling:

```

[edit forwarding-options]
sampling {
 input {
 family inet {
 rate 100;
 run-length 2;
 }
 }
 family inet {
 output {
 file {
 filename sonet-samples.txt;
 files 40;
 size 5m;
 }
 }
 }
}

```

### Example: Sampling All Traffic from a Single IP Address

The following configuration gathers statistical information about every packet entering the router on a specific Gigabit Ethernet port originating from a single source IP address of 172.16.92.31, and collects it in a file named **samples-172-16-92-31.txt**.

Create the filter:

```

[edit firewall family inet]
filter one-ip {
 term get-ip {
 from {
 source-address 172.16.92.31;
 }
 then {
 sample;
 accept;
 }
 }
}

```

Apply the filter to the Gigabit Ethernet interface:

```

[edit interfaces]
ge-4/1/1 {
 unit 0 {
 family inet {
 filter {
 input one-ip;
 }
 address 10.45.92.254;
 }
 }
}

```

```
 }
 }
}
```

Finally, gather statistics on all the candidate samples; in this case, gather all statistics:

```
[edit forwarding-options]
sampling {
 input {
 family inet {
 rate 1;
 }
 }
 family inet {
 output {
 file {
 filename samples-172-16-92-31.txt;
 files 100;
 size 100k;
 }
 }
 }
}
```

---

### Example: Sampling All FTP Traffic

The following configuration gathers statistical information about a moderate percentage of packets using the FTP data transfer protocol in the output path of a specific T3 interface, and collects the information in a file named **t3-ftp-traffic.txt**.

Create a filter:

```
[edit firewall family inet]
filter ftp-stats {
 term ftp-usage {
 from {
 destination-port [ftp ftp-data];
 }
 then {
 sample;
 accept;
 }
 }
}
```

Apply the filter to the T3 interface:

```
[edit interfaces]
t3-7/0/2 {
 unit 0 {
 family inet {
 filter {
 input ftp-stats;
 }
 address 10.35.78.254/32 {
 destination 10.35.78.4;
 }
 }
 }
}
```

```

 }
 }
}

```

Finally, gather statistics on 10 percent of the candidate samples:

```

[edit forwarding-options]
sampling {
 input {
 family inet {
 rate 10;
 }
 }
 family inet {
 output {
 file {
 filename t3-ftp-traffic.txt;
 files 50;
 size 1m;
 }
 }
 }
}

```

- Related Documentation**
- *Traffic Sampling, Forwarding, and Monitoring Overview*
  - [Sampling Instance Configuration on page 881](#)

## Sampling Instance Configuration

You can configure active sampling by defining a sampling instance that specifies a name for the sampling parameters and bind the instance name to an FPC, MPC, or DPC. This configuration enables you to define multiple named sampling parameter sets associated with multiple destinations and protocol families per sampling destination. With the cflowd version 5 and version 8 and flow aggregation version 9, you can use templates to organize the data gathered from sampling.

To implement this feature, you include the **instance** statement at the **[edit forwarding-options sampling]** hierarchy level.

The following considerations apply to the sampling instance configuration:

- This configuration is supported on the IP version 4 (**inet**), IP version 6 (**ipv6**), and MPLS protocol families.
- You can configure the router to perform sampling in either of two locations:
  - On the Routing Engine, using the sampled process. To select this method, use a filter (input or output) with a matching term that contains the **then** sample statement.
  - On the Monitoring Services, Adaptive Services, or Multiservices PIC. Specify the interface name at the **[forwarding-options sampling instance *instance-name* family inet output interface]** hierarchy level. You can configure the same or different services PICs in a set of sampling instances.

- You can configure the **rate** and **run-length** options at the **[edit forwarding-options sampling input]** hierarchy level to apply common values for all families on a global basis. Alternatively, you can configure these options at the **[edit forwarding-options sampling instance *instance-name* input]** hierarchy level to apply specific values for each instance or at the **[edit forwarding-options sampling instance *instance-name* family *family* input]** hierarchy level to apply specific values for each protocol family you configure.
- For MX Series devices with Modular Port Concentrators (MPCs), port-mirrored or sampled packets can be truncated (or clipped) to any length in the range of 1 through 255 bytes. Only the values 1 to 255 are valid for packet truncation on these devices. For other devices, the range is from 0 through 9216. A maximum-packet-length value of zero (0) represents that truncation is disabled, and the entire packet is mirrored or sampled.



**NOTE:** The **run-length** and **maximum-packet-length** configuration statements are not supported on MX80 routers.

---

To associate the defined instance with a particular FPC, MPC, or DPC, you include the **sampling-instance** statement at the **[edit chassis fpc *number*]** hierarchy level, as in the following example:

```
chassis {
 fpc 2 {
 sampling-instance samp1;
 }
}
```

To associate a sampling instance with an FPC in the MX Series Virtual Chassis master or backup router, use the **sampling-instance *instance-name*** statement at the **[edit chassis member *member-number* fpc slot *slot-number*]** hierarchy level, where *member-number* is 0 (for the master router) or 1 (for the backup router), and *slot-number* is a number in the range 0 through 11.

#### Related Documentation

- *Traffic Sampling, Forwarding, and Monitoring Overview*
- *Flow Monitoring Feature Guide for Routing Devices*
- *More Information About Flow Monitoring*
- *Configuring Active Flow Monitoring*
- *Directing Traffic Sampling Output to a Server Running the cflowd Application*
- [Configuring Traffic Sampling on page 871](#)
- *Example: Sampling Instance Configuration*
- *[edit forwarding-options sampling] Hierarchy Level*
- *Inline Flow Monitoring for Virtual Chassis Overview*

## Configuring Discard Accounting

---

Discard accounting is similar to traffic sampling, but varies from it in two ways:

- In discard accounting, the packet is intercepted by the monitoring PIC and is not forwarded to its destination.
- Traffic sampling allows you to limit the number of packets sampled by configuring the **max-packets-per-second**, **rate**, and **run-length** statements. Discard accounting does not provide these options, and a high packet count can potentially overwhelm the monitoring PIC.

A discard instance is a named entity that specifies collector information under the **accounting name** statement. Discard instances are referenced in firewall filter **term** statements by including the **then discard accounting name** statement.

Most of the other statements are also found at the **[edit forwarding-options sampling]** hierarchy level. For information on cflowd, see [“Enabling Flow Aggregation” on page 898](#). The **flow-active-timeout** and **flow-inactive-timeout** statements are described in [“Configuring Flow Monitoring” on page 818](#).

To direct sampled traffic to a flow-monitoring interface, include the **interface** statement. The **engine-id** and **engine-type** statements specify the accounting interface used on the traffic, and the **source-address** statement specifies the traffic source.

You cannot use rate-limiting with discard accounting; however, you can specify the duration of the interval for exporting aggregated accounting information by including the **aggregate-export-interval** statement in the configuration. This enables you to put a boundary on the amount of traffic exported to a flow-monitoring interface.

### Related Documentation

- [Enabling Flow Aggregation on page 898](#)
- [Configuring Flow Monitoring on page 818](#)



# Sampling Data Using Inline Sampling

- [Understanding Inline Active Flow Monitoring on page 885](#)
- [Configuring Inline Active flow Monitoring on page 890](#)
- [Configuring Inline Active Flow Monitoring on MX80 Routers on page 894](#)

## Understanding Inline Active Flow Monitoring

---

This topic provides an overview of the inline active flow monitoring feature and IPFIX and Version 9 flow collection templates used for inline active flow monitoring.

This topic contains the following sections:

- [Inline Active Flow Monitoring on page 885](#)
- [Inline Active Flow Monitoring Limitations and Restrictions on page 886](#)
- [IPFIX and Version 9 Templates on page 887](#)

## Inline Active Flow Monitoring

The inline active flow monitoring is implemented on the Packet Forwarding Engine. All the functions like flow creation, flow update, and flow records export are done by the Packet Forwarding Engine. The flow records are sent out in industry standard IPFIX format.

Inline active flow monitoring provides for higher scalability and performance as the scaling and performance are not dependent on the capacity of the services interface. It is also cost effective in more than one way as there is no need to invest in additional hardware or to dedicate a PIC slot for the services PIC. You can make full use of the available slots for handling traffic on the device.

Junos OS Release 13.2 extends inline active flow monitoring support to VPLS flows. Now, you can configure inline active flow monitoring for IPv4, IPv6, and VPLS traffic.

The inline active flow monitoring configuration can be broadly classified into four categories:

1. Configurations at the **[edit services flow-monitoring]** hierarchy level—At this level, you configure the template properties for inline flow monitoring.
2. Configurations at the **[edit forwarding-options]** hierarchy level—At this level, you configure a sampling instance and associate the template (configured at the **[edit**

**services flow-monitoring**] hierarchy level) with the sampling instance. At this level, you also configure the flow-server IP address and port number as well as the flow export rate.

3. Configurations at the **[edit chassis]** hierarchy level—At this level, you associate the sampling instance with the FPC on which the media interface is present. If you are configuring sampling of IPv6 flows, you must also specify the flow hash table size.
4. Configurations at the **[edit firewall]** hierarchy level—At this level you configure a firewall filter for the family of traffic to be sampled. You must attach this filter to the interface on which you want to sample the traffic.

Inline active flow monitoring supports version 9 and IPFIX flow collection templates. Support for version 9 template was introduced in Junos OS Release 13.2, and is limited to IPv4 flows. IPFIX template is supported for IPv4, IPv6, and VPLS flows. IPFIX template uses UDP as the transport protocol, whereas version 9 is transport protocol-independent.

Before you configure inline active flow monitoring, you should ensure that you have adequately-sized hash tables for IPv4 and IPv6 flow sampling. These tables can use one to fifteen 256k areas, and each table is assigned a default value of one such area. When anticipated traffic volume requires larger tables, allocate larger tables.



**NOTE:** Starting with Junos OS Release 13.3, you can configure flow collectors to be reachable through non-default VPN routing and forwarding (VRF) instances by including the `routing-instance instance-name` statement at the **[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-server hostname]** hierarchy level for inline flow monitoring. You cannot configure a flow collector to be reachable through non-default VRF instances for version 5 and version 8 flows. You must configure the routing instance to be a VRF instance by including the `instance-type vrf` statement at the **[edit routing-instances instance-name]** hierarchy level.

---

## Inline Active Flow Monitoring Limitations and Restrictions

The following limitations and restrictions apply to the inline active flow monitoring feature in Junos OS:

- You can configure inline active flow monitoring only on MX Series routers with Trio-based line cards and T4000 routers with Type 5 FPCs.
- You can apply Version 9 flow template only to IPv4 traffic.
- You can configure only one sampling instance on an Flexible PIC Concentrator (FPC).
- You can configure only one type of sampling—either PIC-based sampling or inline sampling—per family in a sampling instance. However, you can configure PIC-based and inline sampling for different families in a sampling instance.
- You can configure only one collector for inline active flow monitoring.



- The following considerations apply to the inline flow-monitoring instance configuration:
  - Sampling run-length and clip-size are not supported.
  - For inline configurations, each family can support only one collector.
  - The user-defined sampling instance gets precedence over the global instance. When a user-defined sampling instance is attached to the FPC, the global instance is removed from the FPC and the user-defined sampling instance is applied to the FPC.
- On routers with Multiservices PICs or Multiservices DPCs, all fragments of a fragmented IPv4 packet other than the first fragment of the packet are processed accurately by the flow monitoring application running on MS-PIC or MS-DPC. The flow monitoring mechanism handles such fragments accurately by setting the layer 4 related fields in the associated flows to zero.
- Flow records and templates cannot be exported if the flow collector is reachable through any management interface.
- The flow collector should be reachable through the default routing table (inet.0 or inet6.0). If the flow collector is reachable via a non-default VPN routing and forwarding table (VRF), flow records and templates cannot be exported.



**NOTE:** Starting with Junos OS Release 13.3, you can configure the flow collector to be reachable through non-default VRF instances apart from being reachable over the default VRF instance. Flow records and templates can be exported even with non-default VRF instances.

- If the destination of the sampled flow is reachable through multiple paths, the IP\_NEXT\_HOP (Element ID 15) and OUTPUT\_SNMP (Element ID 14) in the IPv4 flow record would be set to the Gateway Address and SNMP Index of the first path seen in the forwarding table.
- If the destination of the sampled flow is reachable through multiple paths, the IP\_NEXT\_HOP (Element ID 15) and OUTPUT\_SNMP (Element ID 14) in the IPv6 flow records would be set to 0.
- The Incoming Interface (IIF) and Outgoing Interface (OIF) should be part of the same VRF. If OIF is in a different VRF, DST\_MASK (Element ID 13), DST\_AS (Element ID 17), IP\_NEXT\_HOP (Element ID 15), and OUTPUT\_SNMP (Element ID 14) would be set to 0 in the flow records.
- Each Lookup Chip (LU) maintains and exports flows independent of other LUs. Traffic received on a media interface is distributed across all LUs in a multi-LU platform. It is likely that a single flow will be processed by multiple LUs. Therefore, each LU creates a unique flow and exports it to the flow collector. This can cause duplicate flows records to be seen on the flow collector. The flow collector should aggregate PKTS\_COUNT and BYTES\_COUNT for duplicate flow records to derive a single flow record.

## IPFIX and Version 9 Templates

The following sections list the fields included in IPFIX and Version 9 templates.

### Fields Included in the IPFIX IPv4 Template

---

- IPv4 Source Address
- IPv4 Destination Address
- IPv4 TOS
- IPv4 Protocol
- L4 Source Port
- L4 Destination Port
- ICMP Type and Code
- Input Interface
- VLAN ID
- IPv4 Source Mask
- IPv4 Destination Mask
- Source AS
- Destination AS
- IPv4 Next Hop Address
- TCP Flags
- Output Interface
- Number of Flow Bytes
- Number of Flow Packets
- Minimum TTL (time to live)
- Maximum TTL (time to live)
- Flow Start Time
- Flow End Time
- Flow End Reason
- 802.1Q VLAN identifier (dot1qVlanId)
- 802.1Q Customer VLAN identifier (dot1qCustomerVlanId)

### Fields Included in the IPFIX IPv6 Template

---

- IPv6 Source Address
- IPv6 Destination Address
- IPv6 TOS
- IPv6 Protocol
- L4 Source Port

- L4 Destination Port
- ICMP Type and Code
- Input Interface
- VLAN ID
- IPv6 Source Mask
- IPv6 Destination Mask
- Source AS
- Destination AS
- IPv6 Next Hop Address
- TCP Flags
- Output Interface
- Number of Flow Bytes
- Number of Flow Packets
- Minimum Hop Limits
- Maximum Hop Limits
- Flow Start Time
- Flow End Time
- Flow End Reason
- 802.1Q VLAN identifier (dot1qVlanId)
- 802.1Q Customer VLAN identifier (dot1qCustomerVlanId)

---

#### Fields Included in the Version 9 IPv4 Template

- IPv4 Source Address
- IPv4 Destination Address
- IPv4 TOS
- IPv4 Protocol
- L4 Source Port
- L4 Destination Port
- ICMP Type and Code
- Input Interface
- VLAN ID
- IPv4 Source Mask
- IPv4 Destination Mask
- Source AS

- Destination AS
- IPv4 Next Hop Address
- BGP IPv4 Next Hop Address
- TCP Flags
- Output Interface
- Number of Flow Bytes
- Number of Flow Packets
- Time when the first packet of the flow was switched.
- Time when the last packet of flow was switched.
- Internet Protocol Version

**Related  
Documentation**

- *Example: Configuring Inline Active Flow Monitoring*
- [Configuring Inline Active Flow Monitoring on MX80 Routers on page 894](#)

---

## Configuring Inline Active flow Monitoring

The inline active flow monitoring is implemented on the Packet Forwarding Engine. All the functions like flow creation, flow update, and flow records export are done by the Packet Forwarding Engine. The flow records are sent out in industry standard IPFIX format.

The inline active flow monitoring configuration can be broadly classified into four categories:

1. Configurations at the **[edit services flow-monitoring]** hierarchy level—At this level, you configure the template properties for inline flow monitoring.
2. Configurations at the **[edit forwarding-options]** hierarchy level—At this level, you configure a sampling instance and associate the template (configured at the **[edit services flow-monitoring]** hierarchy level) with the sampling instance. At this level, you also configure the flow-server IP address and port number as well as the flow export rate.
3. Configurations at the **[edit chassis]** hierarchy level—At this level, you associate the sampling instance with the FPC on which the media interface is present. If you are configuring sampling of IPv6 flows, you must also specify the flow hash table size.
4. Configurations at the **[edit firewall]** hierarchy level—At this level you configure a firewall filter for the family of traffic to be sampled. You must attach this filter to the interface on which you want to sample the traffic.

Before you configure inline active flow monitoring, you should ensure that you have adequately-sized hash tables for IPv4 and IPv6 flow sampling. These tables can use one to fifteen 256k areas, and each table is assigned a default value of one such area. When anticipated traffic volume requires larger tables, allocate larger tables.



**NOTE:** For Junos OS releases earlier than Release 12.1, the following points are applicable for supporting backward compatibility when you configure the IPv4 and IPv6 flow table sizes for inline active flow monitoring:

- If you do not configure the `flow-table-size` statement at the `[edit chassis fpc slot-number inline-services]` hierarchy level, fifteen 256K entries are allocated by default for the IPv4 flow table and one 1K entry is allocated by default for the IPv6 flow table on the Packet Forwarding Engine.
- If you configure the `ipv4-flow-table-size size` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level and if you do not configure the `ipv6-flow-table-size size` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level, the number of units of 256K entries that you configure for the IPv4 flow table is allocated. For the IPv6 flow table, a default size of one 1K entry is allocated on the Packet Forwarding Engine.
- If you do not configure the `ipv4-flow-table-size size` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level and if you configure the `ipv6-flow-table-size size` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level, the number of units of 256K entries that you configure for the IPv6 flow table is allocated. For the IPv4 flow table, a default size of one 1K entry is allocated on the Packet Forwarding Engine.
- If you configure the sizes of both the IPv4 and IPv6 flow tables, the flow tables are created on the Packet Forwarding Engine based on the size that you specified.



**NOTE:** The functionality to log the cflowd records in a log file before they are exported to a cflowd server (by including the `local-dump` statement at the `[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-server hostname]` hierarchy level) is not supported when you configure inline flow monitoring (by including the `inline-jflow` statement at the `[edit forwarding-options sampling instance instance-name family inet output]` hierarchy level).

To allocate IPv4 and IPv6 flow hash tables:

1. Go to the `flow-table-size` hierarchy level for inline services on the FPC that processes the monitored flows.

```
[edit]
user@host# edit chassis fpc 0 inline-services flow-table-size
```

2. Specify the required sizes for the sampling hash tables.

```
[edit chassis fpc 0 inline-services flow-table-size]
user@host# set ipv4-flow-table-size 5
user@host# set ipv6-flow-table-size 5
```



**NOTE:** When you set the flow hash table sizes, remember:

- Any change in the configured size of flow hash table sizes initiates an automatic reboot of the FPC.
- The total number of units used for both IPv4 and IPv6 cannot exceed 15.

To configure inline active flow monitoring on all other MX Series routers (except for MX80 routers), EX Series switches, and T4000 routers with Type 5 FPC:

1. Enable inline active flow monitoring and specify the source address for the traffic.

```
[edit forwarding-options sampling instance instance-name family inet output]
user@host# set inline-jflow source address address
```

2. Specify the IP\_FIX output format.

```
[edit forwarding-options sampling instance instance-name family inet output flow-server
address]
user@host# set version-ipfix template ipv4
```

3. Specify the output properties.

```
[edit services flow-monitoring]
user@host# set version-ipfix
```

The output format properties are common to other output formats and are described in [“Configuring Flow Aggregation to Use IPFIX Flow Templates” on page 912](#).

The following is an example of the sampling configuration for an instance that supports inline active flow monitoring on **family inet** and PIC-based sampling on **family inet6**:

```
[edit forwarding-options]
sampling {
 instance {
 sample-ins1 {
 input {
 rate 1;
 }
 family inet {
 output {
 flow-server 2.2.2.2 {
 port 2055;
 version-ipfix {
 template {
 ipv4;
 }
 }
 }
 }
 inline-jflow {
 source-address 10.11.12.13;
 }
 }
 }
 }
}
```

```

family inet6 {
 output {
 flow-server 2.2.2.2 {
 port 2055;
 version-ipfix {
 template {
 ipv6;
 }
 }
 }
 }
 interface sp-0/1/0 {
 source-address 10.11.12.13;
 }
}
}
}

```

The following example shows the output format configuration:

```

services {
 flow-monitoring {
 version-ipfix {
 template ipv4 {
 flow-active-timeout 60;
 flow-inactive-timeout 60;
 ipv4-template;
 template-refresh-rate {
 packets 1000;
 seconds 10;
 }
 }
 option-refresh-rate {
 packets 1000;
 seconds 10;
 }
 }
 }
}
}

```

The following considerations apply to the inline flow-monitoring instance configuration:

- Sampling run-length and clip-size are not supported.
- For inline configurations, each family can support only one collector.



**NOTE:** On routers with Multiservices PICs or Multiservices DPCs, IPv4 and IPv6 fragments are processed accurately. The flow monitoring application creates two flows for every fragmented flow. The first fragment that has the complete Layer 4 information forms the first flow with 5-tuple data and subsequently, all the fragmented packets related to this flow form another flow with the Layer 4 fields set to zero.

- Related Documentation**
- [Configuring Inline Active Flow Monitoring on MX80 Routers on page 894](#)
  - [inline-jflow on page 1687](#)

---

## Configuring Inline Active Flow Monitoring on MX80 Routers

---

To configure inline active flow monitoring on MX80 routers:

1. Associate a sampling instance with the Forwarding Engine Processor.

```
[edit]
user@host# set chassis tfeb slot number sampling-instance sampling-instance
```

The Forwarding Engine Processor slot is always 0 because MX80 routers have only one Packet Forwarding Engine. In this configuration, the sampling instance is **sample-ins1**.

```
[edit]
user@host# set chassis tfeb slot 0 sampling-instance sample-ins1
```



**NOTE:** MX80 routers support only one sampling instance.

2. Under forwarding-options, configure a sampling instance for the flow server and inline jflow instances (these will be configured in the following steps):

```
[edit forwarding-options sampling]
user@host# edit instance inline_sample
```

3. Configure the rate at the **[edit forwarding-options sampling instance instance-name input]** hierarchy level to apply specific values for the sampling instance **sample-ins1**.

```
[edit forwarding-options sampling instance sample-ins1 input]
user@host# set rate number
```

In this configuration, the rate is 1000.

```
[edit forwarding-options sampling instance sample-ins1 input]
user@host# set rate 1000
```

4. Navigate to the output hierarchy and from there, enable a flow server and then specify the output address and port:

```
[edit] forwarding-options sampling instance inline_sample family inet output]
user@host# edit flow-server address
```

```
[edit forwarding-options sampling instance inline_sample family inet output flow-server
<address>]
user@host# set port number
```

5. Return to the output hierarchy and specify the source address for inline jflow:

```
[edit forwarding-options sampling instance sample-ins1 family inet output]
user@host# set inline-jflow source-address address
```

In this configuration, the source address is 10.11.12.13.

```
[edit forwarding-options sampling instance sample-ins1 family inet output]
```



```
user@host# set inline-jflow source-address 10.11.12.13
```

6. Specify the output properties.

```
[edit services flow-monitoring]
```

```
user@host# set version-ipfix
```

The output format properties are common to other output formats and are described in [“Configuring Flow Aggregation to Use IPFIX Flow Templates” on page 912](#).

The following is an example of the sampling configuration for an instance that supports inline active flow monitoring on MX80 routers:

```
[edit forwarding-options]
user@host# show
sampling {
 instance {
 sample-ins1 {
 input {
 rate 1000;
 }
 family inet {
 flow-server 133..13.13.122{
 port 1333;
 inline-jflow {
 source-address 10.11.12.13;
 }
 }
 }
 }
 }
}
```



**NOTE:** You need not configure a Flexible PIC Concentrator (FPC) slot because MX80 routers have only one Packet Forwarding Engine.

The following considerations apply to the inline flow-monitoring instance configuration:

- This configuration does not support MPLS-IPv6.
- Clip-size is not supported.

#### Related Documentation

- [Configuring Flow Aggregation to Use IPFIX Flow Templates on page 912](#)
- [Configuring Inline Active flow Monitoring on page 890](#)
- [inline-jflow on page 1687](#)



# Sampling Data Using Flow Aggregation

- Understanding Flow Aggregation on page 897
- Enabling Flow Aggregation on page 898
- Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd on page 898
- Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902
- Configuring Flow Aggregation to Use IPFIX Flow Templates on page 912
- Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows on page 918
- Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows on page 921
- Directing Replicated Flows to Multiple Flow Servers on page 926
- Logging cflowd Flows Before Export on page 928

## Understanding Flow Aggregation

---

You can collect an aggregate of sampled flows and send the aggregate to a specified host that runs either the cflowd application available from CAIDA (<http://www.caida.org>) or the newer version 9 format defined in RFC 3954, *Cisco Systems NetFlow Services Export Version 9*. Before you can perform flow aggregation, the routing protocol process must export the autonomous system (AS) path and routing information to the sampling process.

By using flow aggregation, you can obtain various types of byte and packet counts of flows through a router. The application collects the sampled flows over a period of 1 minute. At the end of the minute, the number of samples to be exported are divided over the period of another minute and are exported over the course of the same minute.

You configure flow aggregation in different ways, depending on whether you want to export flow records in cflowd version 5 or 8 format, or the separate version 9 format. The latter allows you to sample MPLS, IPv4, IPv6, and peer AS billing traffic. You can also combine configuration statements between the MPLS and IPv4 formats.



**NOTE:** When PIC-based sampling is enabled, collection of flow statistics for sampled packets on flows in virtual private networks (VPNs) is also supported. No additional CLI configuration is required.

- Related Documentation**
- [Enabling Flow Aggregation on page 898](#)
  - [Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd on page 898](#)
  - [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902](#)
  - [Directing Replicated Flows to Multiple Flow Servers on page 926](#)
  - [Logging cflowd Flows Before Export on page 928](#)

---

## Enabling Flow Aggregation

Before you can perform flow aggregation, the routing protocol process must export the autonomous system (AS) path and routing information to the sampling process. To enable the export of AS path and the routing information to the sampling process, one or more of the following needs to be configured:

- At the **[edit forwarding-options]** hierarchy level (for routing instances, at the **[edit routing-instance *routing-instance-name* forwarding-options]** hierarchy level), configure **sampling family** or **sampling output** or **sampling instance** or **monitoring** or **accounting**.
- At the **[edit routing-options]** hierarchy level (for routing instances, at the **[edit routing-instance *routing-instance-name* routing-options]** hierarchy level), configure **route record**.
- At the **[edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider]** hierarchy level, configure **forwarding-db-size**.

- Related Documentation**
- [Understanding Flow Aggregation on page 897](#)
  - [Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd on page 898](#)
  - [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902](#)
  - [Directing Replicated Flows to Multiple Flow Servers on page 926](#)
  - [Configuring Traffic Sampling on page 871](#)
  - [Example: Configuring Active Flow Monitoring Version 9 for IPv6](#)
  - [Logging cflowd Flows Before Export on page 928](#)

---

## Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd

To enable the collection of cflowd version 5 or version 8 flow formats, include the **flow-server** statement:

```
flow-server hostname {
 aggregation {
 autonomous-system;
 destination-prefix;
 protocol-port;
 source-destination-prefix {
 caida-compliant;
 }
 }
}
```

```

 }
 source-prefix;
 }
 autonomous-system-type (origin | peer);
 (local-dump | no-local-dump);
 port port-number;
 version format;
}

```

You can include this statement at the following hierarchy levels:

- [edit forwarding-options sampling family (inet | inet6 | mpls) output]
- [edit forwarding-options sampling instance *instance-name* output]
- [edit forwarding-options accounting *name* output cflowd *hostname*]

You must configure the **family inet** statement on logical interface **unit 0** on the monitoring interface, as in the following example:

```

[edit interfaces]
sp-3/0/0 {
 unit 0 {
 family inet {
 ...
 }
 }
}

```



**NOTE:** Boot images for monitoring services interfaces are specified at the [edit chassis images pic] hierarchy level. You must enable the NTP client to make the cflowd feature operable, by including the following configuration:

```

[edit system]
ntp {
 boot-server ntp.juniper.net;
 server 172.17.28.5;
}
processes {
 ntp enable;
}

```

For more information, see the *Junos OS Administration Library for Routing Devices*.

You can also configure cflowd version 5 for flow-monitoring applications by including the **cflowd** statement at the [edit forwarding-options monitoring *name* family inet output] hierarchy level:

```

cflowd hostname {
 port port-number;
}

```

The following restrictions apply to cflowd flow formats:

- You can configure up to one version 5 and one version 8 flow format at the **[edit forwarding-options accounting *name* output]** hierarchy level.
- You can configure up to eight version 5 or one version 8 flow format at the **[edit forwarding-options sampling family (inet | inet6 | mpls) output]** hierarchy level for Routing Engine-based sampling by including the **flow-server** statement. In contrast, PIC-based sampling allows you to specify one cflowd version 5 server and one version 8 server simultaneously. However, the two cflowd servers must have different IP addresses.
- You can configure up to eight version 5 flow formats at the **[edit forwarding-options monitoring *name* output]** hierarchy level. Version 8 flow formats and aggregation are not supported for flow-monitoring applications.
- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created on the monitoring PIC only after the route record resynchronization operation is complete, which is 60 seconds after the PIC comes up. Any packets sent to the PIC would be dropped until the synchronization process is complete.
- The configuration includes a proprietary v5 extension template for supporting 4-byte AS information in flow records. Its template version is set to 500, indicating it to be proprietary. All other fields remain the same; the source AS and destination AS are each 4 bytes long, rather than 2 bytes as in the traditional v5 template. This option is available at the **[edit forwarding-options sampling family inet output flow-server server-name version]** hierarchy level.

In the **cflowd** statement, specify the name or identifier of the host that collects the flow aggregates. You must also include the User Datagram Protocol (UDP) port number on the host and the version, which gives the format of the exported cflowd aggregates. To collect cflowd records in a log file before exporting, include the **local-dump** statement.



**NOTE:** You can specify both host (cflowd) sampling and port mirroring in the same configuration; however, only one action takes effect at any one time. Port mirroring takes precedence. For more information, see [“Configuring Port Mirroring” on page 931](#).

For cflowd version 8 only, you can specify aggregation of specific types of traffic by including the **aggregation** statement. This conserves memory and bandwidth by enabling cflowd to export targeted flows rather than all aggregated traffic. To specify a flow type, include the **aggregation** statement:

```
aggregation {
 autonomous-system;
 destination-prefix;
 protocol-port;
 source-destination-prefix {
 caida-compliant;
 }
}
```

```

 }
 source-prefix;
 }

```

You can include this statement at the following hierarchy levels:

- [edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server *hostname*]
- [edit forwarding-options accounting *name* output cflowd *hostname*]

The **autonomous-system** statement configures aggregation by the AS number; this statement might require setting the separate cflowd **autonomous-system-type** statement to include either **origin** or **peer** AS numbers. The **origin** option specifies to use the origin AS of the packet source address in the Source Autonomous System cflowd field. The **peer** option specifies to use the peer AS through which the packet passed in the Source Autonomous System cflowd field. By default, cflowd exports the origin AS number.

The **destination-prefix** statement configures aggregation by the destination prefix only.

The **protocol-port** statement configures aggregation by the protocol and port number; requires setting the separate **cflowd port** statement.

The **source-destination-prefix** statement configures aggregation by the source and destination prefix. Version 2.1b1 of CAIDA's cflowd application does not record source and destination mask length values in compliance with CAIDA's *cflowd Configuration Guide*, dated August 30, 1999. If you configure the **caida-compliant** statement, the Junos OS complies with Version 2.1b1 of cflowd. If you do not include the **caida-compliant** statement in the configuration, the Junos OS records source and destination mask length values in compliance with the *cflowd Configuration Guide*.

The **source-prefix** statement configures aggregation by the source prefix only.

Collection of sampled packets in a local ASCII file is not affected by the **cflowd** statement.

The following commands enable RE- and PIC-based sampling at the **set forwarding options sampling** hierarchy level:

- set input rate *rate*
- set input run-length *length*
- set family inet output flow-server *flowcollector* port *udp port*
- set family inet output flow-server *flowcollector* no-local-dump
- set family inet output flow-server *flowcollector* version <5/8>

The following commands enable RE- and PIC-based sampling at the **set interfaces** hierarchy level:

- *interface to be sampled* unit *unit* family inet filter *input/output filename*

The following commands enable RE- and PIC-based sampling at the **set firewall family** hierarchy level:

- `set inet filter filtername term 1 then count filternameing`
- `set inet filter filtername term 1 then sample`
- `set inet filter filtername term 1 then accept`

The following command enables PIC-based sampling at the **set forwarding options sampling** hierarchy level:

- `set family inet output interface sp-*/*/* source address source address`

The following example shows a PIC-based flow aggregation configuration using version 5:

```
family inet {
 output {
 flow-inactive-timeout 15;
 flow-active-timeout 60;
 flow-server 153.104.248.37 {
 port 9996;
 version 5;
 }
 interface sp-2/2/0 {
 engine-id 4;
 source-address 153.104.0.254;
 }
 }
}
```

The following example shows an RE-based flow aggregation configuration using version 5:

```
family inet {
 output {
 flow-inactive-timeout 15;
 flow-active-timeout 60;
 flow-server 153.104.248.37 {
 port 9996;
 source-address 153.104.0.254;
 version 5;
 }
 }
}
```

#### Related Documentation

- [Understanding Flow Aggregation on page 897](#)
- [Enabling Flow Aggregation on page 898](#)
- [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902](#)
- [Configuring Flow Aggregation to Use IPFIX Flow Templates on page 912](#)

---

## Configuring Flow Aggregation to Use Version 9 Flow Templates

Use of version 9 allows you to define a flow record template suitable for IPv4 traffic, IPv6 traffic, MPLS traffic, a combination of IPv4 and MPLS traffic, or peer AS billing traffic.



Templates and the fields included in the template are transmitted to the collector periodically, and the collector need not be aware of the router configuration.



**NOTE:** Version 9 requires that you install a services PIC, such as the Adaptive Services PIC or Multiservices PIC in the router. On MX Series routers, the Multiservices DPC fulfills this requirement. For more information on determining which services PIC is suitable for your router, see [“Enabling Service Packages” on page 11](#) or the appropriate hardware documentation.



**NOTE:** If multiple protocol families are configured for a particular flow collector, the export packets will originate from multiple Source IDs, with each Source ID corresponding to a particular protocol. The multiple Source IDs do not indicate that the export packets are originating from multiple Service PICs.

The following sections contain additional information:

- [Configuring the Traffic to Be Sampled on page 903](#)
- [Configuring the Version 9 Template Properties on page 904](#)
- [Customizing Template ID, Observation Domain ID, and Source ID for Version 9 flow Templates on page 905](#)
- [Restrictions on page 906](#)
- [Fields Included in Each Template Type on page 906](#)
- [MPLS Sampling Behavior on page 908](#)
- [Verification on page 908](#)
- [Examples: Configuring Version 9 Flow Templates on page 909](#)

## Configuring the Traffic to Be Sampled

To specify sampling of IPv4, IPv6, MPLS, or peer AS billing traffic, include the appropriate configuration of the **family** statement at the **[edit forwarding-options sampling]** hierarchy level:

```
[edit forwarding-options]
sampling {
 family (inet | inet6 | mpls);
}
```

You can include **family inet**, **family inet6**, or **family mpls**.



**NOTE:** If you specify sampling for peer AS billing traffic, the **family** statement supports only IPv4 and IPv6 traffic (inet or inet6). Peer AS billing traffic is enabled only at the global instance hierarchy level and is not available for per Packet Forwarding Engine instances.

After you specify the family of traffic to be sampled, configure the sampling parameters such as the maximum packet length (beyond which the packets are truncated), maximum packets to be sampled per second (beyond which the packets are dropped), the rate (for example, if you specify 10, every 10th packet is sampled), and run length (which specify the number of packets to be sampled after the trigger; that is if the **rate** is set to 10 and **run-length** to 5, five packets starting the 10th packet are sampled).

```
[edit forwarding-options sampling]
input {
 maximum-packet-length bytes
 max-packets-per-second number;
 rate number;
 run-length number;
}
```

## Configuring the Version 9 Template Properties

To define the version 9 templates, include the following statements at the **[edit services flow-monitoring version9]** hierarchy level:

```
[edit services flow-monitoring version9]
template name {
 options-template-id
 template-id
 source-id
 flow-active-timeout seconds;
 flow-inactive-timeout seconds;
 option-refresh-rate packets packets seconds seconds;
 template-refresh-rate packets packets seconds seconds;
 (ipv4-template | ipv6-template | mpls-ipv4-template | mpls-template |
 peer-as-billing-template) {
 label-position [positions];
 }
}
```

The following details apply to the configuration statements:

- You assign each template a unique name by including the **template name** statement.
- You then specify each template for the appropriate type of traffic by including the **ipv4-template**, **ipv6-template**, **mpls-ipv4-template**, **mpls-template**, or **peer-as-billing-template**.
- If the template is used for MPLS traffic, you can also specify up to three label positions for the MPLS header label data by including the **label-position** statement; the default values are **[1 2 3]**.
- Within the template definition, you can optionally include values for the **flow-active-timeout** and **flow-inactive-timeout** statements. These statements have specific default and range values when they are used in template definitions; the default is 60 seconds and the range is from 10 through 600 seconds. Values you specify in template definitions override the global timeout values configured at the **[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server]** hierarchy level.

- You can also include settings for the **option-refresh-rate** and **template-refresh-rate** statements within a template definition. For both of these properties, you can include a timer value (in seconds) or a packet count (in number of packets). For the **seconds** option, the default value is 60 and the range is from 10 through 600. For the **packets** option, the default value is 4800 and the range is from 1 through 480,000.
- To filter IPV6 traffic on a media interface, the following configuration is supported:

```

interfaces interface-name {
 unit 0 {
 family inet6 {
 sampling {
 input;
 output;
 }
 }
 }
}

```

## Customizing Template ID, Observation Domain ID, and Source ID for Version 9 flow Templates

Use of version 9 and IPFIX allows you to define a flow record template suitable for IPv4 traffic, IPv6 traffic, MPLS traffic, a combination of IPv4 and MPLS traffic, or peer AS billing traffic. Templates and the fields included in the template are transmitted to the collector periodically, and the collector need not be aware of the router configuration. Starting with Junos OS Release 14.1, you can specify the unique identifier for the version 9 and IPFIX templates. The identifier of a template is locally unique within a combination of a transport session and an observation domain. Template IDs 0 through 255 are reserved for template sets, options template sets, and other sets for future use. Template IDs of data sets are numbered from 256 through 65535. Typically, this information element or field in the template is used to define the characteristics or properties of other information elements in a template. After a restart of the export process of templates is performed, template IDs can be reassigned. In Junos OS releases earlier than Release 14.1, template IDs and options template IDs were predefined for each address family and could not be modified.

This functionality to configure template ID, options template ID, observation domain ID, and source ID is supported on all routers with MPCs (Trio chip-based FPCs).

The following values were assigned by default for the template IDs of IPFIX templates for the different protocols or address families, until Junos OS Release 13.3:

- IPv4 flow template ID—256
- IPv6 flow template ID—257
- VPLS flow template ID—258
- Options template ID for all address families—512

The corresponding data sets and option data sets contain the value of the template IDs and options template IDs respectively in the set ID field. This method enables the collector to match a data record with a template record.

For more information about specifying the source ID, observation domain ID, template ID, and options template ID for version 9 and IPFIX flows, see [“Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows” on page 918](#) and [“Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows” on page 921](#).

## Restrictions

The following restrictions apply to version 9 templates:

- You cannot apply the two different types of flow aggregation configuration (cflowd version 5/8 and flow aggregation version 9) at the same time.
- Flow export based on an **mpls-ipv4** template assumes that the IPv4 header follows the MPLS header. In the case of Layer 2 VPNs, the packet on the provider router (P router) would look like this:

MPLS | Layer 2 Header | IPv4

In this case, **mpls-ipv4** flows are not created on the PIC, because the IPv4 header does not directly follow the MPLS header. Packets are dropped on the PIC and are accounted as parser errors.

- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created on the monitoring PIC only after the route record resynchronization operation is complete, which is 60 seconds after the PIC comes up. Any packets sent to the PIC would be dropped until the synchronization process is complete.



**NOTE:** "Because the forwarding of a packet that arrives with MPLS labels is performed based on the MPLS label and not based on the IP address contained in the packet, the packet is sampled at the output interface with the MPLS label that was popped not being available at the time of sampling. In such a case, depending on the incoming interface (IIF), the VRF index is identified and the route for the sampled packet is determined in the VRF table. Because a specific route is not available in the VRF that is different from the VRF on which the packet is received, the Output Interface Index, Source Mask, and Destination Mask fields are incorrectly populated. This behavior occurs when an IPv4 template is applied as a firewall filter on an egress interface with sample as the action."

---

## Fields Included in Each Template Type

The following fields are common to all template types:

- Input interface
- Output interface
- Number of bytes

- Number of packets
- Flow start time
- Flow end time

The IPv4 template includes the following specific fields:

- IPv4 Source Address
- IPv4 Destination Address
- L4 Source Port
- L4 Destination Port
- IPv4 TOS
- IPv4 Protocol
- ICMP type and code
- TCP Flags
- IPv4 Next Hop Address
- Source autonomous system (AS) number
- Destination AS number

The IPv6 template includes the following specific fields:

- IPv6 Source Address and Mask
- IPv6 Destination Address and Mask
- L4 Source Port
- L4 Destination Port
- IPv6 TOS
- IPv6 Protocol
- TCP Flags
- IP Protocol Version
- IPv6 Next Hop Address
- Egress Interface Information
- Source Autonomous System (AS) number
- Destination AS number

The MPLS template includes the following specific fields:

- MPLS Label #1
- MPLS Label #2
- MPLS Label #3

- MPLS EXP Information
- FEC IP Address

The MPLS-IPv4 template includes all the fields found in the IPv4 and MPLS templates.

The peer AS billing template includes the following specific fields:

- IPv4 Class of Service (TOS)
- Ingress Interface
- BGP IPv4 Next Hop Address
- BGP Peer Destination AS Number

## MPLS Sampling Behavior

This section describes the behavior when MPLS sampling is used on egress interfaces in various scenarios (label pop or swap) on provider routers (P routers). For more information on configuration and background specific to MPLS applications, see the *MPLS Applications Feature Guide for Routing Devices*.

1. You configure MPLS sampling on an egress interface on the P router and configure an MPLS flow aggregation template. The route action is label *pop* because penultimate hop popping (PHP) is enabled.

Previously, IPv4 packets (only) would have been sent to the PIC for sampling even though you configured MPLS sampling. No flows should be created, with the result that the parser fails.

With the current capability of applying MPLS templates, MPLS flows are created.

2. As in the first case, you configure MPLS sampling on an egress interface on the P router and configure an MPLS flow aggregation template. The route action is label *swap* and the swapped label is 0 (explicit null).

The resulting behavior is that MPLS packets are sent to the PIC. The flow being sampled corresponds to the label before the swap.

3. You configure a Layer 3 VPN network, in which a customer edge router (CE-1) sends traffic to a provider edge router (PE-A), through the P router, to a similar provider edge router (PE-B) and customer edge router (CE-2) on the remote end.

The resulting behavior is that you cannot sample MPLS packets on the PE-A to P router link.

## Verification

To verify the configuration properties, you can use the **show services accounting aggregation template template-name name** operational mode command.

All other **show services accounting** commands also support version 9 templates, except for **show services accounting flow-detail** and **show services accounting aggregation aggregation-type**. For more information about operational mode commands, see the [CLI Explorer](#).

## Examples: Configuring Version 9 Flow Templates

The following is a sample version 9 template configuration:

```
services {
 flow-monitoring {
 version9 {
 template ip-template {
 flow-active-timeout 20;
 flow-inactive-timeout 120;
 ipv4-template;
 }
 template mpls-template-1 {
 mpls-template {
 label-position [1 3 4];
 }
 }
 template mpls-ipv4-template-1 {
 mpls-ipv4-template {
 label-position [1 5 7];
 }
 }
 template peer-as-billing-template-1 {
 peer-as-billing-template;
 }
 }
 }
}
```

The following is a sample firewall filter configuration for MPLS traffic:

```
firewall {
 family mpls {
 filter mpls_sample {
 term default {
 then {
 accept;
 sample;
 }
 }
 }
 }
}
```

The following sample configuration applies the MPLS sampling filter on a networking interface and configures the AS PIC to accept both IPv4 and MPLS traffic:

```
interfaces {
 at-0/1/1 {
 unit 0 {
 family mpls {
 filter {
 input mpls_sample;
 }
 }
 }
 }
}
```

```
 }
 }
 sp-7/0/0 {
 unit 0 {
 family inet;
 family mpls;
 }
 }
}
```

The following example applies the MPLS version 9 template to the sampling output and sends it to the AS PIC:

```
forwarding-options {
 sampling {
 input {
 family mpls {
 rate 1;
 }
 }
 family mpls {
 output {
 flow-active-timeout 60;
 flow-inactive-timeout 30;
 flow-server 1.2.3.4 {
 port 2055;
 version9 {
 template mpls-ipv4-template-1;
 }
 }
 }
 }
 interface sp-7/0/0 {
 source-address 1.1.1.1;
 }
 }
}
```

The following is a sample firewall filter configuration for the peer AS billing traffic:

```
firewall {
 family inet {
 filter peer-as-filter {
 term 0 {
 from {
 destination-class dcu-1;
 interface ge-2/1/0;
 forwarding-class class-1;
 }
 then count count_team_0;
 }
 }
 term 1 {
 from {
 destination-class dcu-2;
 interface ge-2/1/0;
 }
 }
 }
}
```



```

 forwarding-class class-1;
 }
 then count count_team_1;
}
term 2 {
 from {
 destination-class dcu-3;
 interface ge-2/1/0;
 forwarding-class class-1;
 }
 then count count_team_2;
}
}
}
}

```

The following sample configuration applies the peer AS firewall filter as a filter attribute under the forwarding-options hierarchy for CoS-level data traffic usage information collection:

```

forwarding-options {
 family inet {
 filter output peer-as-filter;
 }
}

```

The following sample configuration applies the peer AS DCU policy options to collect usage statistics for the traffic stream for as-path ingressing at a specific input interface with the firewall configuration hierarchy applied as Forwarding Table Filters (FTFs). The configuration functionality with COS capability can be achieved through FTFs for destination-class usage with forwarding-class for specific input interfaces:

```

policy-options {
 policy-statement P1 {
 from {
 protocol bgp;
 neighbor 10.2.25.5; #BGP router configuration;
 as-path AS-1; #AS path configuration;
 }
 then destination-class dcu-1; #Destination class configuration;
 }
 policy-statement P2 {
 from {
 neighbor 1.2.25.5;
 as-path AS-2;
 }
 then destination-class dcu2;
 }
 policy-statement P3 {
 from {
 protocol bgp;
 neighbor 192.2.1.1;
 as-path AS-3;
 }
 then destination-class dcu3;
 }
}

```

```
as-path AS-1 3131:1111:1123;
as-path AS-2 100000;
as-path AS-3 192:29283:2;
}
```

The following example applies the peer-as-billing version 9 template to enable sampling of traffic for billing purposes:

```
forwarding-options {
 sampling {
 }
 input {
 rate 1;
 }
 family inet {
 output {
 flow-server 10.209.15.58 {
 port 300;
 version9 {
 template {
 peer-as;
 }
 }
 }
 }
 interface sp-5/2/0 {
 source-address 2.3.4.5;
 }
 }
}
}
family inet {
 filter {
 output peer-as-filter;
 }
}
```

**Related  
Documentation**

- [Understanding Flow Aggregation on page 897](#)
- [Enabling Flow Aggregation on page 898](#)
- [Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd on page 898](#)
- [Configuring Flow Aggregation to Use IPFIX Flow Templates on page 912](#)
- [Configuring Traffic Sampling on page 871](#)
- *Example: Configuring Active Flow Monitoring Version 9 for IPv6*

---

## Configuring Flow Aggregation to Use IPFIX Flow Templates

Use of IPFIX allows you to define a flow record template suitable for IPv4 traffic or IPv6 traffic. Templates are transmitted to the collector periodically, and the collector need not be aware of the router configuration. You can define template refresh rate, flow active timeout and inactive timeout.

If flow records are being sent for multiple protocol families (for example, for IPv4 and IPv6), each protocol family flow will have a unique Observation Domain ID.

The following sections contain additional information:

- [Configuring the IPFIX Template Properties on page 913](#)
- [Restrictions on page 914](#)
- [Customizing Template ID, Observation Domain ID, and Source ID for IPFIX flow Templates on page 914](#)
- [Fields Included in the IPv4 Template on page 915](#)
- [Fields Included in the IPv6 Template on page 916](#)
- [Verification on page 916](#)
- [Example: Configuring an IPFIX Flow Templates and Flow Sampling on page 917](#)

## Configuring the IPFIX Template Properties

To define the IPFIX templates, include the following statements at the **[edit services flow-monitoring version-ipfix]** hierarchy level:

```
[edit services flow-monitoring IPFIX]
template name {
 options-template-id
 template-id
 observation-domain-id
 flow-active-timeout seconds;
 flow-inactive-timeout seconds;
 option-refresh-rate packets packets seconds seconds;
 template-refresh-rate packets packets seconds seconds;
 (ipv4-template | ipv6-template);
}
```

The following details apply to the configuration statements:

- You assign each template a unique name by including the **template *name*** statement.
- You then specify each template for the appropriate type of traffic by including the **ipv4-template** or **ipv6-template**.
- Within the template definition, you can optionally include values for the **flow-active-timeout** and **flow-inactive-timeout** statements. These statements have specific default and range values when they are used in template definitions; the default is 60 seconds and the range is from 10 through 600 seconds.
- You can also include settings for the **option-refresh-rate** and **template-refresh-rate** statements within a template definition. For both of these properties, you can include a timer value (in seconds) or a packet count (in number of packets). For the **seconds** option, the default value is 600 and the range is from 10 through 600. For the **packets** option, the default value is 4800 and the range is from 1 through 480,000.
- To filter IPV6 traffic on a media interface, the following configuration is supported:

```
interfaces interface-name {
 unit 0 {
```

```
family inet6 {
 sampling {
 input;
 output;
 }
}
```

## Restrictions

The following restrictions apply to IPFIX templates:

- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created only after the route record resynchronization operation is complete, which takes 120 seconds.
- VLAN ID field is not valid for egress traffic, and returns a value of 0 for egress traffic.
- The VLAN ID field is updated when a new flow record is created and so, any change in VLAN ID after the record has been created might not be updated in the record.

## Customizing Template ID, Observation Domain ID, and Source ID for IPFIX flow Templates

Use of version 9 and IPFIX allows you to define a flow record template suitable for IPv4 traffic, IPv6 traffic, MPLS traffic, a combination of IPv4 and MPLS traffic, or peer AS billing traffic. Templates and the fields included in the template are transmitted to the collector periodically, and the collector need not be aware of the router configuration. Starting with Junos OS Release 14.1, you can specify the unique identifier for the version 9 and IPFIX templates. The identifier of a template is locally unique within a combination of a transport session and an observation domain. Template IDs 0 through 255 are reserved for template sets, options template sets, and other sets for future use. Template IDs of data sets are numbered from 256 through 65535. Typically, this information element or field in the template is used to define the characteristics or properties of other information elements in a template. After a restart of the export process of templates is performed, template IDs can be reassigned. In Junos OS releases earlier than Release 14.1, template IDs and options template IDs were predefined for each address family and could not be modified.

This functionality to configure template ID, options template ID, observation domain ID, and source ID is supported on all routers with MPCs (Trio chip-based FPCs).

The following values were assigned by default for the template IDs of version 9 templates for the different protocols or address families, until Junos OS Release 13.3:

- IPv4 flow template ID—272
- IPv6 flow template ID—273

- VPLS flow template ID—274
- Options template ID for all address families—520

The corresponding data sets and option data sets contain the value of the template IDs and options template IDs respectively in the set ID field. This method enables the collector to match a data record with a template record.

For more information about specifying the source ID, observation domain ID, template ID, and options template ID for version 9 and IPFIX flows, see [“Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows” on page 918](#) and [“Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows” on page 921](#).

### Fields Included in the IPv4 Template

- IPv4 Source Address
- IPv4 Destination Address
- IPv4 TOS
- IPv4 Protocol
- L4 Source Port
- L4 Destination Port
- ICMP Type and Code
- Input Interface
- VLAN ID
- IPv4 Source Mask
- IPv4 Destination Mask
- Source AS
- Destination AS
- IPv4 Next Hop Address
- TCP Flags
- Output Interface
- Number of Flow Bytes
- Number of Flow Packets
- Minimum TTL (time to live)
- Maximum TTL (time to live)
- Flow Start Time
- Flow End Time
- Flow End Reason

- 802.1Q VLAN identifier (dot1qVlanId)
- 802.1Q Customer VLAN identifier (dot1qCustomerVlanId)

### Fields Included in the IPv6 Template

- IPv6 Source Address
- IPv6 Destination Address
- IPv6 TOS
- IPv6 Protocol
- L4 Source Port
- L4 Destination Port
- ICMP Type and Code
- Input Interface
- VLAN ID
- IPv6 Source Mask
- IPv6 Destination Mask
- Source AS
- Destination AS
- IPv6 Next Hop Address
- TCP Flags
- Output Interface
- Number of Flow Bytes
- Number of Flow Packets
- Minimum Hop Limits
- Maximum Hop Limits
- Flow Start Time
- Flow End Time
- Flow End Reason
- 802.1Q VLAN identifier (dot1qVlanId)
- 802.1Q Customer VLAN identifier (dot1qCustomerVlanId)
- Fragment Identification
- IPv6 Extension Headers

### Verification

The following show commands are supported for IPFIX:

- `show services accounting flow inline-jflow fpc-slot fpc-slot`
- `show services accounting errors inline-jflow fpc-slot fpc-slot`
- `show services accounting status inline-jflow fpc-slot fpc-slot`

### Example: Configuring an IPFIX Flow Templates and Flow Sampling

The following is a sample IPFIX template configuration:

```
services {
 flow-monitoring {
 version-ipfix {
 template ipv4 {
 flow-active-timeout 60;
 flow-inactive-timeout 70;
 template-refresh-rate seconds 30;
 option-refresh-rate seconds 30;
 ipv4-template;
 }
 }
 }
}

chassis {
 fpc 0 {
 sampling-instance s1;
 }
}
```

The following example applies the IPFIX template to enable sampling of traffic for billing:

```
forwarding-options {
 sampling {
 instance {
 s1 {
 input {
 rate 10;
 }
 family inet {
 output {
 flow-server 11.11.4.2 {
 port 2055;
 version-ipfix {
 template {
 ipv4;
 }
 }
 }
 }
 inline-jflow {
 source-address 11.11.2.1;
 }
 }
 }
 }
 }
}
```

- Related Documentation**
- [Understanding Flow Aggregation on page 897](#)
  - [Inclusion of Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates](#)
  - [Enabling Flow Aggregation on page 898](#)
  - [Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd on page 898](#)
  - [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902](#)

## Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows

---

For IPFIX flows, an identifier of an Observation Domain is locally unique to an exporting process of the templates. The export process uses the Observation Domain ID to uniquely identify to the collection process in which the flows were metered. We recommend that you configure this ID to be unique for each IPFIX flow. A value of 0 indicates that no specific Observation Domain is identified by this information element. Typically, this attribute is used to limit the scope of other information elements. If the observation domain is not unique, the collector cannot uniquely identify an IPFIX device.

If you configure the same Observation Domain ID for different template types, such as for IPv4 and IPv6, it does not impact flow monitoring because the actual or the base observation domain ID is transmitted in the flow. The actual observation domain ID is derived from the value you configure and also in conjunction with other parameters such as the slot number, lookup chip (LU) instance, Packet Forwarding Engine instance. Such a method of computation of the observation domain ID ensures that this ID is not the same for two IPFIX devices.

Until Junos OS Release 13.3, the observation domain ID is predefined and is set to a fixed value, which is derived from the combination of FPC slot, sampling protocol, PFE Instance and LU Instance fields. This derivation creates a unique observation domain per LU per family. Starting with Junos OS Release 14.1, you can configure the observation domain ID, which causes the first 8 bits of the field to be configured.

The following modifications have been made:

- FPC slots are expanded to 8 bits to enable more slots to be configured in an MX Series Virtual Chassis configuration.
- 8 bits of the configured observation domain ID are used.
- You can configure a value for the observation domain ID in the range of 0 through 255.
- The Protocol field is increased to 3 bits to provide support for additional protocols in inline flow monitoring.
- You can associate the observation domain ID with templates by using the **observation-domain-id *domain-id*** statement at the **[edit services flow- monitoring version-ipfix template *template-name*]** hierarchy level.

For version 9 flows, a 32-bit value that identifies the Exporter Observation Domain is called the source ID. NetFlow collectors use the combination of the source IP address



and the source ID field to separate different export streams originating from the same exporter.

To specify the observation domain ID for IPFIX flows, include the **observation-domain-id domain-id** statement at the **[edit services flow-monitoring version-ipfix template template-name]** hierarchy level.

```
[edit services flow-monitoring version-ipfix]
template template-name {
 observation-domain-id domain-id;
}
```

To specify the source ID for version 9 flows, include the **source-id source-id** statement at the **[edit services flow-monitoring version9 template template-name]** hierarchy level.

```
[edit services flow-monitoring version9]
template template-name {
 source-id source-id;
}
```

Table 32 on page 919 describes observation domain ID values for different combinations of the configured domain ID, protocol family, FPC slot, and the Packet Forwarding Engine and lookup chip instances.

**Table 32: Example of Observation Domain ID**

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id Conf val rsvd 1proto slot LUInst PFEInst  xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
None	IPv4 (0)	1	1	0	0000 0000 0000 1000 0000 0001 0000 0001 0x00080101
None	IPv6 (1)	1	1	0	0000 0000 0000 1001 0000 0001 0000 0001 0x00090101
None	VPLS (2)	1	1	0	0000 0000 0000 1010 0000 0001 0000 0001 0x000A0101

Table 32: Example of Observation Domain ID *(continued)*

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id  Conf val rsvd lproto slot LUInst PFEInst  xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
None	MPLS (3)	1	1	0	0000 0000 0000 1011 0000 0001 0000 0001 0x000B0101
4	IPv4 (0)	1	1	0	0000 0100 0000 1000 0000 0001 0000 0001 0x04080101
190	IPv4 (0)	1	1	0	1101 1110 0000 1000 0000 0001 0000 0001 0xBE080101
4	IPv4 (0)	2	1	1	0000 0100 0000 1000 0000 0010 0001 0001 0x04080211
4	IPv6 (1)	1	1	0	0000 0100 0000 1001 0000 0001 0001 0000 0x04090110
190	IPv6 (1)	1	1	0	1101 1110 0000 1001 0000 0001 0001 0000 0xBE090110
4	VPLS (2)	2	2	0	0000 0100 0000 1010 0000 0010 0010 0000 0x040A0220

Table 32: Example of Observation Domain ID (*continued*)

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id  Conf val rsvd lproto slot LUInst PFEInst  xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
10	IPv4 (0)	28	2	1	0000 1010 0000 1000 0001 1100 0010 0001 0x0A081C21

**Related Documentation** • [Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows on page 921](#)

## Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows

Starting with Junos OS Release 14.1, you can define the template ID for version 9 and IPFIX templates for inline flow monitoring. To specify the template ID for version 9 flows, include the **template-id** *id* statement at the **[edit services flow-monitoring version9 template *template-name*]** hierarchy level.

```
[edit services flow-monitoring version9]
template template-name {
 template-id id;
}
```

To specify the template ID for version IPFIX flows, include the **template-id** statement at the **[edit services flow-monitoring version-ipfix template *template-name*]** hierarchy level.

```
[edit services flow-monitoring version-ipfix]
template template-name {
 template-id id;
}
```

To specify the options template ID for version 9 flows, include the **options-template-id** statement at the **[edit services flow-monitoring version9 template *template-name*]** hierarchy level.

```
[edit services flow-monitoring version9]
template template-name {
 options-template-id id;
}
```

To specify the options template ID for version IPFIX flows, include the **options-template-id** statement at the **[edit services flow-monitoring version-ipfix template *template-name*]** hierarchy level. The template ID and options template ID can be a value in the range of 1024 through 65535.

```
[edit services flow-monitoring version-ipfix]
template template-name {
 options-template-id id;
}
```

The template ID and options template ID can be a value in the range of 1024 through 65535. If you do not configure values for the template ID and options template ID, default values are assumed for these IDs, which are different for the various address families. If you configure the same template ID or options template ID value for different address families, such a setting is not processed properly and might cause unexpected behavior. For example, if you configure the same template ID value for both IPv4 and IPv6, the collector validates the export data based on the template ID value that it last receives. In this case, if IPv6 is configured after IPv4, the value is effective for IPv6 and the default value is used for IPv4.

The following are the default values of template IDs for IPFIX flows for the different protocols or address families, until Junos OS Release 13.3:

- IPv4 IPFIX flow template ID—256
- IPv6 IPFIX flow template ID—257
- VPLS IPFIX flow template ID—258
- MPLS IPFIX flow template ID—259

The following are the default values of template IDs for version 9 flows for the different protocols or address families, starting with Junos OS Release 14.1:

- IPv4 version 9 flow template ID—320
- IPv6 version 9 flow template ID—321
- VPLS version 9 flow template ID—322
- MPLS version 9 flow template ID—323

The following are the default values of template IDs for IPFIX flows for the different protocols or address families, until Junos OS Release 13.3:

- IPv4 IPFIX flow options template ID—512
- IPv6 IPFIX flow options template ID—513
- VPLS IPFIX flow options template ID—514
- MPLS IPFIX flow options template ID—515

The following are the default values of template IDs for version 9 flows for the different protocols or address families, starting with Junos OS Release 14.1:

- IPv4 version 9 flow options template ID—576
- IPv6 version 9 flow options template ID—577
- VPLS version 9 flow options template ID—578
- MPLS version 9 flow options template ID—579

[Table 33 on page 923](#) describes the values of data template and option template IDs for different protocols with default and configured values for IPFIX flows.

**Table 33: Values of Template and Option Template IDs for IPFIX Flows**

Family	Configured Value	Data Template	Option Template
IPv4	None	256	576
IPv4	1024-65535	1024-65535	1024-65535
IPv6	None	257	577
IPv6	1024-65535	1024-65535	1024-65535
VPLS	None	258	578
VPLS	1024-65535	1024-65535	1024-65535
MPLS	None	259	579
MPLS	1024-65535	1024-65535	1024-65535

[Table 34 on page 923](#) describes the values of data template and option template IDs for different protocols with default and configured values for version 0 flows.

**Table 34: Values of Template and Option Template IDs for Version 9 Flows**

Family	Configured Value	Data Template	Option Template
IPv4	None	320	576
IPv4	1024-65535	1024-65535	1024-65535
IPv6	None	321	577
IPv6	1024-65535	1024-65535	1024-65535
VPLS	None	322	578
VPLS	1024-65535	1024-65535	1024-65535

**Table 34: Values of Template and Option Template IDs for Version 9 Flows (continued)**

Family	Configured Value	Data Template	Option Template
MPLS	None	323	579
MPLS	1024-65535	1024-65535	1024-65535

Table 33 on page 923 describes the values of data template and option template IDs for different protocols with default and configured values for IPFIX flows.

**Table 35: Values of Template and Option Template IDs for IPFIX Flows**

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id Conf val rsvd 1proto slot LUInst PFEInst xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
None	IPv4 (0)	1	1	0	0000 0000 0000 1000 0000 0001 0000 0001 0x00080101
None	IPv6 (1)	1	1	0	0000 0000 0000 1001 0000 0001 0000 0001 0x00090101
None	VPLS (2)	1	1	0	0000 0000 0000 1010 0000 0001 0000 0001 0x000A0101
None	MPLS (3)	1	1	0	0000 0000 0000 1011 0000 0001 0000 0001 0x000B0101
4	IPv4 (0)	1	1	0	0000 0100 0000 1000 0000 0001 0000 0001 0x04080101

**Table 35: Values of Template and Option Template IDs for IPFIX Flows (continued)**

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id Conf val rsvd 1proto slot LUInst PFEInst  xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
190	IPv4 (0)	1	1	0	1101 1110 0000 1000 0000 0001 0000 0001 0xBE080101
4	IPv4 (0)	2	1	1	0000 0100 0000 1000 0000 0010 0001 0001 0x04080211
4	IPv6 (1)	1	1	0	0000 0100 0000 1001 0000 0001 0001 0000 0x04090110
190	IPv6 (1)	1	1	0	1101 1110 0000 1001 0000 0001 0001 0000 0xBE090110
4	VPLS (2)	2	2	0	0000 0100 0000 1010 0000 0010 0010 0000 0x040A0220
10	IPv4 (0)	28	2	1	0000 1010 0000 1000 0001 1100 0010 0001 0x0A081C21

**Related Documentation**

- [Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows on page 918](#)

## Directing Replicated Flows to Multiple Flow Servers

You can configure replication of the sampled flow records for use by multiple flow servers. You can use either sampling based on the Routing Engine, using cflowd version 5 or version 8, or sampling based on the services PIC, using flow aggregation version 9, as described in the following sections:

- [Directing Replicated Routing Engine–Based Sampling Flows to Multiple Servers on page 926](#)
- [Directing Replicated Version 9 Flow Aggregates to Multiple Servers on page 927](#)

### Directing Replicated Routing Engine–Based Sampling Flows to Multiple Servers

Routing Engine–based sampling supports up to eight flow servers for both cflowd version 5 and version 8 configurations. The total number of servers is limited to eight regardless of how many are configured for cflowd v5 or v8.

When you configure cflowd-based sampling, the export packets are replicated to all flow servers configured to receive them. If two servers are configured to receive v5 records, both the servers will receive records for a specified flow.



**NOTE:** With Routing Engine–based sampling, if multiple flow servers are configured with version 8 export format, all of them must use the same aggregation type. For example, all servers receiving version 8 export could be configured for source-destination aggregation type.

The following configuration example allows replication of export packets to two flow servers.

```
forwarding-options {
 sampling {
 instance inst1 {
 input {
 rate 1;
 }
 family inet;
 output {
 flow-server 10.10.3.2 {
 port 2055;
 version 5;
 source-address 192.168.164.119;
 }
 flow-server 172.17.20.62 {
 port 2055;
 version 5;
 source-address 192.168.164.119;
 }
 }
 }
 }
}
```



```
 }
 }
}
```

## Directing Replicated Version 9 Flow Aggregates to Multiple Servers

The export packets generated for a template are replicated to all the flow servers that are configured to receive information for that template. The maximum number of servers supported is eight.

This also implies that periodic updates required by version 9 (RFC 3954) are sent to each configured collector. The following updates are sent periodically as part of this requirement:

- Options data
- Template definition

The refresh period for options data and template definition is configured on a per-template basis at the **[edit services flow-monitoring]** hierarchy level.

The following configuration example allows replication of version 9 export packets to two flow servers.

```
forwarding-options {
 sampling {
 instance inst1 {
 input {
 rate 1;
 }
 family inet;
 output {
 flow-server 10.10.3.2 {
 port 2055;
 version9 {
 template {
 ipv4;
 }
 }
 }
 flow-server 172.17.20.62 {
 port 2055;
 version9 {
 template {
 ipv4;
 }
 }
 }
 }
 flow-inactive-timeout 30;
 flow-active-timeout 60;
 interface sp-4/0/0 {
 source-address 10.10.3.4;
 }
 }
 }
}
```

```
}
}
```

**Related Documentation**

- [Active Flow Monitoring Overview on page 815](#)
- [Configuring Flow Monitoring on page 818](#)
- [Configuring Services Interface Redundancy with Flow Monitoring on page 826](#)
- [Example: Configuring Active Monitoring on Logical Systems on page 823](#)

---

## Logging cflowd Flows Before Export

To collect the cflowd flows in a log file before they are exported, include the **local-dump** statement at the **[edit forwarding-options sampling output flow-server *hostname*]** hierarchy level:

```
[edit forwarding-options sampling output flow-server hostname]
local-dump;
```

By default, the flows are collected in **/var/log/sampled**; to change the filename, include the **filename** statement at the **[edit forwarding-options sampling traceoptions]** hierarchy level. For more information about changing the filename, see [“Configuring Traffic Sampling Output” on page 876](#).



**NOTE:** Because the **local-dump** statement adds extra overhead, you should use it only while debugging cflowd problems, not during normal operation.

The following is an example of the flow information. The AS number exported is the origin AS number. All flows that belong under a cflowd header are dumped, followed by the header itself:

```
Jun 27 18:35:43 v5 flow entry
Jun 27 18:35:43 Src addr: 192.53.127.1
Jun 27 18:35:43 Dst addr: 192.6.255.15
Jun 27 18:35:43 Nhop addr: 192.6.255.240
Jun 27 18:35:43 Input interface: 5
Jun 27 18:35:43 Output interface: 3
Jun 27 18:35:43 Pkts in flow: 15
Jun 27 18:35:43 Bytes in flow: 600
Jun 27 18:35:43 Start time of flow: 7230
Jun 27 18:35:43 End time of flow: 7271
Jun 27 18:35:43 Src port: 26629
Jun 27 18:35:43 Dst port: 179
Jun 27 18:35:43 TCP flags: 0x10
Jun 27 18:35:43 IP proto num: 6
Jun 27 18:35:43 TOS: 0xc0
Jun 27 18:35:43 Src AS: 7018
Jun 27 18:35:43 Dst AS: 11111
Jun 27 18:35:43 Src netmask len: 16
Jun 27 18:35:43 Dst netmask len: 0
```

[... 41 more version 5 flow entries; then the following header:]

```
Jun 27 18:35:43 cflowd header:
Jun 27 18:35:43 Num-records: 42
Jun 27 18:35:43 Version: 5
Jun 27 18:35:43 low seq num: 118
Jun 27 18:35:43 Engine id: 0
Jun 27 18:35:43 Engine type: 3
```

**Related  
Documentation**

- [Active Flow Monitoring Overview on page 815](#)
- [Configuring Flow Monitoring on page 818](#)
- [Directing Replicated Flows to Multiple Flow Servers on page 926](#)
- [Configuring Services Interface Redundancy with Flow Monitoring on page 826](#)
- [Example: Configuring Active Monitoring on Logical Systems on page 823](#)



## CHAPTER 63

# Sending Packets for Analysis Using Port Mirroring

- [Understanding Port Mirroring on page 931](#)
- [Configuring Port Mirroring on page 931](#)
- [Defining a Next-Hop Group for Port Mirroring on page 948](#)
- [Example: Multiple Port Mirroring with Next-Hop Groups Configuration on page 949](#)

## Understanding Port Mirroring

---

On routers containing an Internet Processor II application-specific integrated circuit (ASIC) or T Series Internet Processor, you can send a copy of an IP version 4 (IPv4) or IP version 6 (IPv6) packet from the router to an external host address or a packet analyzer for analysis. This is known as *port mirroring*.

Port mirroring is different from traffic sampling. In traffic sampling, a sampling key based on the IPv4 header is sent to the Routing Engine. There, the key can be placed in a file, or cflowd packets based on the key can be sent to a cflowd server. In port mirroring, the entire packet is copied and sent out through a next-hop interface.

You can configure simultaneous use of sampling and port mirroring, and set an independent sampling rate and run-length for port-mirrored packets. However, if a packet is selected for both sampling and port mirroring, only one action can be performed and port mirroring takes precedence. For example, if you configure an interface to sample every packet input to the interface and a filter also selects the packet to be port mirrored to another interface, only the port mirroring would take effect. All other packets not matching the explicit filter port-mirroring criteria continue to be sampled when forwarded to their final destination.

### Related Documentation

- [Configuring Port Mirroring on page 931](#)
- [Example: Multiple Port Mirroring with Next-Hop Groups Configuration on page 949](#)

## Configuring Port Mirroring

---

To prepare traffic for port mirroring, include the **filter** statement at the **[edit firewall family inet]** hierarchy level:

```
filter filter-name;
```

This filter at the `[edit firewall family (inet | inet6)]` hierarchy level selects traffic to be port-mirrored:

```
filter filter-name {
 term term-name {
 then {
 port-mirror;
 accept;
 }
 }
}
```

To configure port mirroring on a logical interface, configure the following statements at the `[edit forwarding-options port-mirroring]` hierarchy level:

```
[edit forwarding-options port-mirroring family inet]
input {
 maximum-packet-length bytes;
 rate rate;
 run-length number;
}
family (inet|inet6) {
 output {
 interface interface-name {
 next-hop address;
 }
 no-filter-check;
 }
}
```

or

```
[edit forwarding-options port-mirroring]
input {
 maximum-packet-length bytes;
 rate rate;
 run-length number;
}
family inet6 {
 output {
 next-hop-group group-name{
 group-type inet6;
 interface interface-name {
 next-hop ipv6-address;
 }
 }
 next-hop-subgroup group-name{
 interface interface-name {
 next-hop ipv6-address;
 }
 }
 }
}
```



**NOTE:** The input statement can also be configured at the [edit forwarding-options port-mirroring] hierarchy level. This is only maintained for backward compatibility. However, the configuration of the output statement is deprecated at the [edit forwarding-options port-mirroring] hierarchy level.

Specify the port-mirroring destination by including the **next-hop** statement at the [edit forwarding-options port-mirroring output interface *interface-name*] hierarchy level:

```
next-hop address;
```



**NOTE:** For IPv4 port mirroring to reach a next-hop destination, you must manually include a static Address Resolution Protocol (ARP) entry in the router configuration.

You can also specify the port-mirroring destination by including the **next-hop-group** statement at the [edit forwarding-options port-mirroring family inet6 output] hierarchy level:

```
next-hop-group group-name {
 group-type inet6;
 interface interface-name {
 next-hop ipv6-address;
 }
 next-hop-subgroup group-name {
 interface interface-name {
 next-hop ipv6-address;
 }
 }
}
```

The **no-filter-check** statement is required when you send port-mirrored traffic to a Tunnel PIC that has a filter applied to it. en

The interface used to send the packets to the analyzer is the output interface configured above at the [edit forwarding-options port-mirroring family (inet | inet6) output] hierarchy level. You can use any physical interface type, including generic routing encapsulation (GRE) tunnel interfaces. The next-hop address specifies the destination address; this statement is mandatory for non point-to-point interfaces, such as Ethernet interfaces.

To configure the sampling rate or duration, include the **rate** or **run-length** statement at the [edit forwarding-options port-mirroring input] hierarchy level.

You can trace port-mirroring operations the same way you trace sampling operations. For more information, see [“Tracing Traffic Sampling Operations” on page 878](#).

For more information about port mirroring, see the following sections:

- [Configuring Tunnels on page 934](#)
- [Port Mirroring with Next-Hop Groups on page 936](#)

- [Configuring Inline Port Mirroring on page 937](#)
- [Filter-Based Forwarding with Multiple Monitoring Interfaces on page 938](#)
- [Restrictions on page 938](#)
- [Configuring Port Mirroring on Services Interfaces on page 939](#)
- [Examples: Configuring Port Mirroring on page 940](#)

## Configuring Tunnels

In typical applications, you send the sampled packets to an analyzer or a workstation for analysis, rather than another router. If you must send this traffic over a network, you should use tunnels. For more information about tunnel interfaces, see *Tunnel Properties*.

The MX Series routers support Dense Port Concentrators (DPCs) with built-in Ethernet ports, which do not support Tunnel Services PICs. To create tunnel interfaces on an MX Series router with a DPC, you configure the DPC and the corresponding Packet Forwarding Engine to use for tunneling services at the **[edit chassis]** hierarchy level. You also configure the amount of bandwidth reserved for tunnel services. The Junos OS creates tunnel interfaces on the Packet Forwarding Engine.

To create tunnel interfaces on MX Series routers, include the following statements at the **[edit chassis]** hierarchy level:

```
[edit chassis]
fpc slot-number {
 pic number {
 tunnel-services {
 bandwidth bandwidth-value;
 }
 }
}
```

Include the **fpc slot-number** statement to specify the slot number of the DPC. If two SCBs are installed, the range is 0 through 1. If three SCBs are installed, the range is 0 through 5 and 7 through 11.

Include the **pic number** statement to specify the number of the Packet Forwarding Engine on the DPC. The range is 0 through 3.

You can also specify the amount of bandwidth to allocate for tunnel traffic on each Packet Forwarding Engine by including the **bandwidth bandwidth-value** statement at the **[edit chassis fpc slot-number pic pic-number]** hierarchy level:

- **1g** indicates that 1 Gbps of bandwidth is reserved for tunnel traffic. Configure this option only for a Packet Forwarding Engine on a Gigabit Ethernet 40-port DPC.
- **10g** indicates that 10 Gbps of bandwidth is reserved for tunnel traffic. Configure this option only for a Packet Forwarding Engine on a 10-Gigabit Ethernet 4-port DPC.
- **20g** or **40g**—Configure 20 gigabits per second or 40 gigabits per second only on an MX Series router with the MPC3E and the 100-Gigabit CFP MIC.

If you specify a bandwidth that is not compatible with the type of DPC and Packet Forwarding Engine, tunnel services are not activated. For example, you cannot specify a



bandwidth of 1 Gbps for a Packet Forwarding Engine on a 10-Gigabit Ethernet 4-port DPC.

When you configure tunnel interfaces on the Packet Forwarding Engine of a 10-Gigabit Ethernet 4-port DPC, the Ethernet interfaces for that port are removed from service and are no longer visible in the command-line interface (CLI). The Packet Forwarding Engine of a 10-Gigabit Ethernet 4-port DPC supports either tunnel interfaces or Ethernet interfaces, but not both. Each port on the 10-Gigabit Ethernet 4-port DPC includes two LEDs, one for tunnel services and one for Ethernet services, to indicate which type of service is being used. On the Gigabit Ethernet 40-port DPC, you can configure both tunnel and Ethernet interfaces at the same time.

If your router is equipped with a Tunnel PIC, you can forward duplicate packets to multiple interfaces by configuring a next-hop group. To configure a next-hop group, include the **next-hop-group** statement at the **[edit forwarding-options]** hierarchy level:

```
[edit forwarding-options]
next-hop-group group-names {
 interface interface-name {
 next-hop address;
 }
}
```

The **interface** statement specifies the interface that sends out sampled information. The **next-hop** statement specifies the next-hop addresses to which to send the sampled information.

For IPv6 port mirroring to reach next-hop destination, you can configure a **next-hop-group** statement at the **[edit forwarding-options port-mirroring family inet6 output]** hierarchy level:

```
next-hop-group group-name {
 group-type inet6;
 interface interface-name {
 next-hop ipv6-address;
 }
 next-hop-subgroup group-name {
 interface interface-name {
 next-hop ipv6-address;
 }
 }
}
```

Next-hop groups have the following restrictions:

- Next-hop groups are supported for inet, inet6, and bridge family.
- Next-hop groups are supported on M Series and MX Series routers.
- Next-hop groups or next-hop subgroups support up to 16 next-hop addresses.
- Up to 30 next-hop groups are supported.
- Each next-hop group is expected to have at least two next-hop addresses.
- Each next-hop subgroup supports up to 16 next-hop groups.

## Port Mirroring with Next-Hop Groups

You can configure next-hop groups for M Series, MX Series, and TX Series routers using either IP addresses or Layer 2 addresses for the next hops. Use the **group-type [ inet | inet6 | layer-2 ]** statement at **[edit forwarding-options next-hop-group next-hop-group-name]** hierarchy level to establish the next-hop groups. You can reference more than one port mirroring instance in a filter on MX Series routers. Use the **port-mirror-instance instance-name** statement at the **[edit firewall family family-name filter filter-name term term-name]** hierarchy level to refer to one of several port mirroring instances. For more information about this configuration, see the *Layer 2 Port Mirroring Feature Guide for Routing Devices*.



**NOTE:** On MX Series routers with MPCs, port mirroring instances can only be bound to the FPC level and not up to the PIC level. For MX series routers with a DPC card, both levels are supported.

On M Series, MX Series, and T Series routers only, you can configure port mirroring using next-hop groups, also known as *multipacket port mirroring*, without the presence of a Tunnel PIC. To configure this functionality, include the **next-hop-group** statement at the **[edit forwarding-options port-mirror family [inet | inet6] output]** or **[edit forwarding-options port-mirror instance instance-name family inet output]** hierarchy level:

```
[edit forwarding-options]
port-mirror {
 family inet {
 output {
 next-hop-group group-name {
 interface interface-name {
 next-hop address;
 }
 }
 }
 }
}
or
[edit forwarding-options]
port-mirror {
 family inet6 {
 output {
 next-hop-group group-name {
 group-type inet6;
 interface interface-name {
 next-hop ipv6-address;
 }
 }
 next-hop-subgroup group-name {
 interface interface-name {
 next-hop ipv6-address;
 }
 }
 }
 }
}
```

```

 }
 }
}
or
[edit forwarding-options]
port-mirror {
 instance instance-name {
 family (inet | vpls) {
 output {
 next-hop-group group-name;
 }
 }
 }
}

```

You define the next-hop group by including the **next-hop-group** statement at the **[edit forwarding-options]** hierarchy level. For an example, see [“Examples: Configuring Port Mirroring” on page 940](#). This configuration is supported with IPv4 and IPv6 addresses.

You can disable this configuration by including a **disable** or **disable-all-instances** statement at the **[edit forwarding-options port-mirror]** hierarchy level or by including a **disable** statement at the **[edit forwarding-options port-mirror instance *instance-name*]** hierarchy level. You can display the settings and network status by issuing the **show forwarding-options next-hop-group** and **show forwarding-options port-mirroring** operational commands.



**NOTE:** If you try to bind any derived instance to the FPC, a commit error will occur.

## Configuring Inline Port Mirroring

Inline port mirroring provides you with the ability to specify instances that are not bound to the flexible PIC concentrator (FPC) in the firewall filter’s **then port-mirror-instance** action. This way, you are not limited to only two port-mirror instances per FPC. Inline port mirroring decouples the port-mirror destination from the input parameters like **rate**. While the input parameters are programmed in the switch interface board, the next-hop destination of the mirrored packet is available in the packet itself. Inline port mirroring is supported only on MX Series routers with MPCs.

Using inline port mirroring, a port-mirror instance will have an option to inherit input parameters from another instance that specifies it, as shown in the following CLI configuration example:

```

instance pm2 {
 + input-parameters-instance pm1;
 family inet {
 output {
 interface ge-1/2/3.0 {
 next-hop 50.0.0.3;
 }
 }
 }
}

```

```
 }
 }
}
```

Multiple levels of inheritance are not allowed. One instance can be referred by multiple instances. An instance can refer to another instance that is defined before it. Forward references are not allowed and an instance cannot refer to itself, doing so will cause an error during configuration parsing.

The user can specify an instance that is not bound to the FPC in the firewall filter. The specified filter should inherit one of the two instances that have been bound to the FPC. If it does not, the packet is not marked for port-mirroring. If it does, then the packet will be sampled using the input parameters specified by the referred instance but the copy will be sent to the its own destination.

## Filter-Based Forwarding with Multiple Monitoring Interfaces

If port-mirrored packets are to be distributed to multiple monitoring or collection interfaces based on patterns in packet headers, it is helpful to configure a filter-based forwarding (FBF) filter on the port-mirroring egress interface.

When an FBF filter is installed as an output filter, a packet that is forwarded to the filter has already undergone at least one route lookup. After the packet is classified at the egress interface by the FBF filter, it is redirected to another routing table for additional route lookup. Obviously, the route lookup in the latter routing table (designated by an FBF routing instance) must result in a different next hop from those from the previous tables the packet has passed through, to avoid packet looping inside the Packet Forwarding Engine.

For more information about FBF configuration, see the *Junos OS Routing Protocols Library for Routing Devices*. For an example of FBF applied to an output interface, see [“Examples: Configuring Port Mirroring” on page 940](#).

## Restrictions

The following restrictions apply to port-mirroring configurations:

- The interface you configure for port mirroring should not participate in any kind of routing activity.
- The destination address you specify should not have a route to the ultimate traffic destination. For example, if the sampled IPv4 packets have a destination address of **10.68.9.10** and the port-mirrored traffic is sent to **10.68.20.15** for analysis, the device associated with the latter address should not know a route to **10.68.9.10**. Also, it should not send the sampled packets back to the source address.
- IPv4 and IPv6 traffic is supported. For IPv6 port mirroring, you must configure the next-hop router with an IPv6 neighbor before mirroring the traffic, similar to an ARP request for IPv4 traffic. All the restrictions applied to IPv4 configurations should also apply to IPv6.
- On M120 and M320 routers, multiple next-hop mirroring is not supported.

- Because M320 routers do not support multiple bindings of port-mirror instances per FPC, the **port-mirror-instance** action is not supported in firewall filters for these routers.
- Port mirroring in the ingress and egress direction is not supported for link services IQ (lsq-) interfaces.
- On M Series routers other than the M120 and M320 routers, only one family protocol (either IPv4 or IPv6) is supported at a time.
- Port mirroring supports up to 16 next hops.
- Only transit data is supported.
- You can configure multiple port-mirroring interfaces per router.
- On routers containing an Internet Processor II application-specific integrated circuit (ASIC), you must include a firewall filter with both the **accept** action and the **port-mirror** action modifier on the inbound interface. Do not include the **discard** action, or port mirroring will not work.
- If the port-mirroring interface is a non-point-to-point interface, you must include an IP address under the **port-mirroring** statement to identify the other end of the link. This IP address must be reachable for you to see the sampled traffic. If the port-mirroring interface is an Ethernet interface, the router should have an Address Resolution Protocol (ARP) entry for it. The following sample configuration sets up a static ARP entry.
- You do not need to configure firewall filters on both inbound and outbound interfaces, but at least one is necessary on the inbound interface to provide the copies of the packets to send to an analyzer.
- Inline port mirroring is supported only on MX Series routers with MPCs.
- Configuration for both port mirroring and traffic sampling are handled by the same daemon, so in order to view a trace log file for port mirroring, you must configure the **traceoptions** option under traffic sampling.

## Configuring Port Mirroring on Services Interfaces

A special situation arises when you configure unit **0** of a services interface (AS or Multiservices PIC) to be the port-mirroring logical interface, as in the following example:

```
[edit forwarding-options]
port-mirroring {
 input {
 rate 1;
 }
 family inet {
 output {
 interface sp-1/0/0.0;
 }
 }
}
```

Since any traffic directed to unit **0** on a services interface is targeted for monitoring (cflowd packets are generated for it), the sample port-mirroring configuration indicates

that the customer would like to have cflowd records generated for the port-mirrored traffic.

However, generation of cflowd records requires the following additional configuration; if it is missing, the port-mirrored traffic is simply dropped by the services interface without generating any cflowd packets.

```
[edit forwarding-options]
sampling {
 instance instance1 { # named instances of sampling parameters
 input {
 rate 1;
 }
 family inet {
 output {
 flow-server 172.16.28.65 {
 port 1230;
 }
 }
 interface sp-1/0/0 { # If the port-mirrored traffic requires monitoring, this
 # interface must be same as that specified in the
 # port-mirroring configuration.
 source-address 3.1.2.3;
 }
 }
 }
}
```



**NOTE:** Another way to configure sp-1/0/0 to generate cflowd records is to use only the sampling configuration, but include a firewall filter `sample` action instead of a `port-mirror` action.

---

## Examples: Configuring Port Mirroring

The following example sends port-mirrored traffic to multiple cflowd servers or packet analyzers:

```
[edit interfaces]
ge-1/0/0 { # This is the input interface where packets enter the router.
 unit 0 {
 family inet {
 filter {
 input mirror_pkts; # Here is where you apply the first filter.
 }
 address 10.11.0.1/24;
 }
 }
}
ge-1/1/0 { # This is an exit interface for HTTP packets.
 unit 0 {
 family inet {
 address 10.12.0.1/24;
 }
 }
}
```

```

 }
 }
 ge-1/2/0 { # This is an exit interface for HTTP packets.
 unit 0 {
 family inet {
 address 10.13.0.1/24;
 }
 }
 }
 so-0/3/0 { # This is an exit interface for FTP packets.
 unit 0 {
 family inet {
 address 10.1.1.1/30;
 }
 }
 }
 so-4/3/0 { # This is an exit interface for FTP packets.
 unit 0 {
 family inet {
 address 10.2.2.2/30;
 }
 }
 }
 so-7/0/0 { # This is an exit interface for all remaining packets.
 unit 0 {
 family inet {
 address 10.5.5.5/30;
 }
 }
 }
 so-7/0/1 { # This is an exit interface for all remaining packets.
 unit 0 {
 family inet {
 address 10.6.6.6/30;
 }
 }
 }
 vt-3/3/0 { # The tunnel interface is where you send the port mirrored traffic.
 unit 0 {
 family inet;
 }
 unit 1 {
 family inet {
 filter {
 input collect_pkts; # This is where you apply the second firewall filter.
 }
 }
 }
 }
}
[edit forwarding-options]
port-mirroring { # This is required when you configure next-hop groups.
 input {
 rate 1; # This rate port mirrors one packet for every one received (1:1 = all
 # packets).
 }
 family inet {

```

```
 output { # This sends traffic to a tunnel interface to prepare for multiport mirroring.
 interface vt-3/3/0.1;
 no-filter-check;
 }
 }
}
next-hop-group ftp-traffic { # Point-to-point interfaces require you to specify the interface
 # name only.
 interface so-4/3/0.0;
 interface so-0/3/0.0;
}
next-hop-group http-traffic { # You need to configure a next hop for multipoint interfaces
 # (Ethernet).
 interface ge-1/1/0.0 {
 next-hop 10.12.0.2;
 }
 interface ge-1/2/0.0 {
 next-hop 10.13.0.2;
 }
}
next-hop-group default-collect {
 interface so-7/0/0.0;
 interface so-7/0/1.0;
}
[edit firewall]
family inet {
 filter mirror_pkts { # Apply this filter to the input interface.
 term catch_all {
 then {
 count input_mirror_pkts;
 port-mirror; # This action sends traffic to be copied and port mirrored.
 accept;
 }
 }
 }
 filter collect_pkts { # Apply this filter to the tunnel interface.
 term ftp-term { # This term sends FTP traffic to an FTP next-hop group.
 from {
 protocol ftp;
 }
 then next-hop-group ftp-traffic;
 }
 term http-term { # This term sends HTTP traffic to an HTTP next-hop group.
 from {
 protocol http;
 }
 then next-hop-group http-traffic;
 }
 term default { # This term sends all remaining traffic to a final next-hop group.
 then next-hop-group default-collectors;
 }
 }
}
```



The following example demonstrates configuration of filter-based forwarding at the output interface. In this example, the packet flow follows this path:

1. A packet arrives at interface **fe-1/2/0.0** with source and destination addresses **10.50.200.1** and **10.50.100.1**, respectively.
2. The route lookup in routing table **inet.0** points to the egress interface **so-0/0/3.0**.
3. The output filter installed at **so-0/0/3.0** redirects the packet to routing table **fbf.inet.0**.
4. The packet matches the entry **10.50.100.0/25**, and finally leaves the router from interface **so-2/0/0.0**.

```
[edit interfaces]
so-0/0/3 {
 unit 0 {
 family inet {
 filter {
 output fbf;
 }
 address 10.50.10.2/25;
 }
 }
}
fe-1/2/0 {
 unit 0 {
 family inet {
 address 10.50.50.2/25;
 }
 }
}
so-2/0/0 {
 unit 0 {
 family inet {
 address 10.50.20.2/25;
 }
 }
}
[edit firewall]
filter fbf {
 term 0 {
 from {
 source-address {
 10.50.200.0/25;
 }
 }
 then routing-instance fbf;
 }
 term d {
 then count d;
 }
}
[edit routing-instances]
fbf {
 instance-type forwarding;
 routing-options {
```

```
 static {
 route 10.50.100.0/25 next-hop so-2/0/0.0;
 }
 }
}
[edit routing-options]
interface-routes {
 rib-group inet fbf-group;
}
static {
 route 10.50.100.0/25 next-hop 10.50.10.1;
}
rib-groups {
 fbf-group {
 import-rib [inet.0 fbf.inet.0];
 }
}
```

The following example shows configuration of port mirroring using next-hop groups or multipacket port mirroring:

```
forwarding-options {
 next-hop-group inet_nhg {
 group-type inet;
 interface ge-2/0/2.101 {
 next-hop 10.2.0.2;
 }
 interface ge-2/2/8.2 {
 next-hop 10.8.0.2;
 }
 }
 next-hop-group vpls_nhg {
 group-type layer-2;
 interface ge-2/0/1.100;
 interface ge-2/2/9.0;
 inactive: next-hop-subgroup vpls_subg {
 interface ge-2/0/1.101;
 interface ge-2/2/9.1;
 }
 }
 next-hop-group vpls_nhg_2 {
 group-type layer-2;
 interface ge-2/2/1.100;
 interface ge-2/3/9.0;
 }
}
port-mirror {
 disable-all-instances; /* Disable all port-mirroring instances */
 disable; /* Disable the global instance */
 input {
 rate 10; # start mirroring every 10th packet
 run-length 4; # mirror 4 additional packets
 }
 family inet {
 output {
 next-hop-group inet_nhg;
 }
 }
}
```

```

}
family inet6 {
 output {
 next-hop-group inet6_nhg6 {
 group-type inet6;
 interface ge-2/0/3.102 {
 next-hop 10::1:1:10;
 }
 interface ge-2/0/4.103 {
 next-hop 20::1:1:10;
 }
 next-hop-subgroup vpls_subg {
 interface ge-2/0/.101 {
 next-hop 3::1:1:1;
 }
 interface ge-2/2/9.1 {
 next-hop 4::1:1:1;
 }
 }
 }
 }
}
family vpls {
 output {
 next-hop-group vpls_nhg;
 }
}
instance {
 inst1 {
 disable; /* Disable this instance */
 input {
 rate 1;
 maximum-packet-length 200;
 }
 family inet {
 output {
 next-hop-group inet_nhg;
 }
 }
 family inet6 {
 output {
 next-hop-group inet6_nhg6;
 }
 }
 family vpls {
 output {
 next-hop-group vpls_nhg_2;
 }
 }
 }
}
}

```

The following example shows configuration of port mirroring using next-hop groups or multipacket port mirroring on a T Series router:

```
forwarding-options {
 next-hop-group inet_nhg {
 group-type inet;
 interface so-0/0/0.0; # There is no need for the nexthop address on T Series routers
 interface ge-2/0/2.0 {
 next-hop 1.2.3.4
 }
 }
 next-hop-subgroup sub_inet {
 interface so-1/2/0.0;
 interface ge-6/1/2.0 {
 next-hop 6.7.8.9;
 }
 }
 next-hop-group vpls_nhg_2 {
 group-type layer-2;
 interface ge-2/2/1.100;
 interface ge-2/3/9.0;
 }
}
port-mirroring {
 disable-all-instances; /*Disable all port-mirroring instances */
 disable; /* Disable the global instance */
 input {
 rate 10;
 run-length 4;
 }
 family inet {
 output {
 next-hop-group inet_nhg;
 }
 }
 family vpls {
 output {
 next-hop-group vpls_nhg;
 }
 }
 instance {
 inst1 {
 disable; /* Disable this instance */
 input {
 rate 1;
 maximum-packet-length 200;
 }
 family inet {
 output {
 next-hop-group inet_nhg;
 }
 }
 family vpls {
 output {
 next-hop-group vpls_nhg_2;
 }
 }
 }
 }
}
```

```
}

```

The following example shows configuration of inline port mirroring using PM1, PM2, and PM3 as our port mirror instances.

```
instance {
 pm1 {
 input {
 rate 3;
 }
 family inet {
 output {
 interface ge-1/2/2.0 {
 next-hop 40.0.0.2;
 }
 }
 }
 }
 pm2 {
 input-parameters-instance pm1;
 family inet {
 output {
 interface ge-1/2/3.0 {
 next-hop 50.0.0.3;
 }
 }
 }
 }
 pm3 {
 input {
 rate 3;
 }
 family inet6 {
 output {
 interface ge-1/2/3.0 {
 next-hop 5::5:5:1;
 }
 }
 }
 }
}
firewall {
 filter pm_filter {
 term t1 {
 then port-mirror-instance pm2;
 }
 }
 filter nhg6_filter6 {
 term t6 {
 then next-hop-group inet6-nhg6;
 }
 }
}
chassis {
 fpc 1 {
 port-mirror-instance pm1;
 }
}
```

```
}
```

The packets will be sampled at a rate of 3, and the copy is sent to 50.0.0.3.

**Related  
Documentation**

- [Understanding Port Mirroring on page 931](#)
- [Example: Multiple Port Mirroring with Next-Hop Groups Configuration on page 949](#)

---

## Defining a Next-Hop Group for Port Mirroring

On routers containing an Internet Processor II application-specific integrated circuit (ASIC) or T Series Internet Processor, you can send a copy of an IP version 4 (IPv4) or IP version 6 (IPv6) packet from the router to an external host address or a packet analyzer for analysis. This is known as *port mirroring*.

Port mirroring is different from traffic sampling. In traffic sampling, a sampling key based on the IPv4 header is sent to the Routing Engine. There, the key can be placed in a file, or cflowd packets based on the key can be sent to a cflowd server. In port mirroring, the entire packet is copied and sent out through a next-hop interface.

You can configure simultaneous use of sampling and port mirroring, and set an independent sampling rate and run-length for port-mirrored packets. However, if a packet is selected for both sampling and port mirroring, only one action can be performed, and port mirroring takes precedence. For example, if you configure an interface to sample every packet input to the interface and a filter also selects the packet to be port mirrored to another interface, only the port mirroring would take effect. All other packets not matching the explicit filter port-mirroring criteria continue to be sampled when forwarded to their final destination.

Next-hop groups allow you to include port mirroring multiple interfaces used to forward duplicate packets used in port mirroring.

On MX Series routers, you can mirror tunnel interface input traffic to multiple destinations. To this form of multipacket port mirroring, you specify two or more additional destinations in a next-hop group, define a firewall filter that references the next-hop group as the filter action, and then apply the filter to a logical tunnel interface (lt-) or virtual tunnel interface (vt-) on the MX Series router.

To define a next-hop group for a Layer 2 port-mirroring firewall filter action:

1. Enable the configuration of forwarding options.

```
[edit]
user@host set forwarding-options port-mirroring family (inet | inet6) output
```

2. Enable configuration of a next-hop-group for Layer 2 port mirroring.

```
[edit forwarding-options port-mirroring ... family (inet | inet6) output]
user@host# set next-hop-group next-hop-group-name
```

3. Specify the type of addresses to be used in the next-hop group configuration.

```
[edit forwarding-options port-mirroring ... family (inet | inet6) output next-hop-group
next-hop-group-name]
user@host# set group-type inet6
```

- Specify the interfaces of the next-hop route.

```
[edit forwarding-options port-mirroring ... family (inet | inet6) output next-hop-group
next-hop-group-name]
user@host# set interface logical-interface-name-1
user@host# set interface logical-interface-name-2
```

or

```
[edit forwarding-options port-mirroring ... family (inet | inet6) output next-hop-group
next-hop-group-name]
user@host# set interface interface-name next-hop next-hop-address
```

The MX Series router supports up to 30 next-hop groups. Each next-hop group supports up to 16 next-hop addresses. Each next-hop group must specify at least two addresses. The *next-hop-address* can be an IPv4 or IPv6 address.

- (Optional) Specify the next-hop subgroup.

```
[edit forwarding-options port-mirroring ... family (inet | inet6) output next-hop-group
next-hop-group-name]
user@host# set next-hop-subgroup subgroup-name interface interface-name next-hop
next-hop-address
```

- Verify the configuration of the next-hop group.

```
[edit forwarding-options port-mirroring ... family (inet | inet6) output next-hop-group
next-hop-group-name]
user@host# top
[edit]
user@host# show forwarding-options
```

```
...
next-hop-group next-hop-group-name {
 group-type inet6;
 interface logical-interface-name-1;
 interface interface-name{
 next-hop next-hop-address;
 }
 next-hop-subgroup subgroup-name{
 interface interface-name{
 next-hop next-hop-address;
 }
 }
}
...
```

#### Related Documentation

- [Configuring Port Mirroring on page 931](#)
- [Example: Multiple Port Mirroring with Next-Hop Groups Configuration on page 949](#)
- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*

## Example: Multiple Port Mirroring with Next-Hop Groups Configuration

When you need to analyze traffic containing more than one packet type, or you wish to perform multiple types of analysis on a single type of traffic, you can implement multiple

port mirroring and next-hop groups. You can make up to 16 copies of traffic per group and send the traffic to next-hop group members. A maximum of 30 groups can be configured on a router at any given time. The port-mirrored traffic can be sent to any interface, except aggregated SONET/SDH, aggregated Ethernet, loopback (**lo0**), or administrative (**fxp0**) interfaces. To send port-mirrored traffic to multiple flow servers or packet analyzers, you can use the **next-hop-group** statement at the **[edit forwarding-options]** hierarchy level.

**Figure 35: Active Flow Monitoring—Multiple Port Mirroring with Next-Hop Groups Topology Diagram**

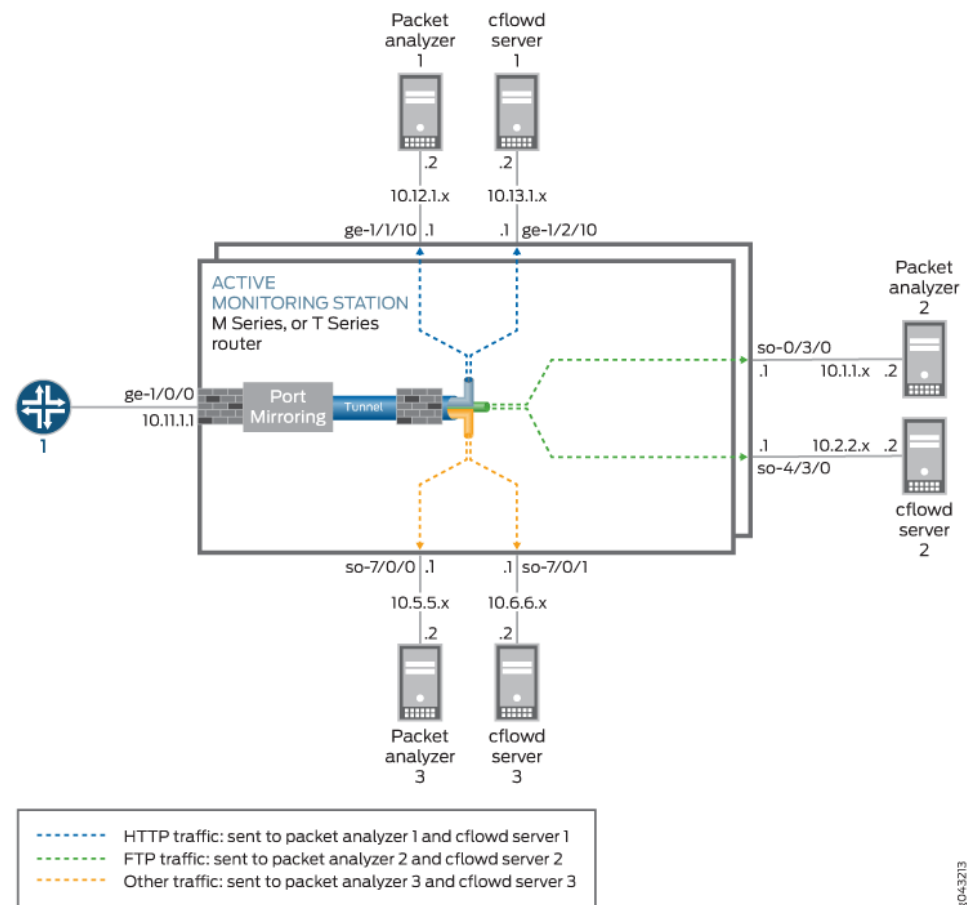


Figure 35 on page 950 shows an example of how to configure multiple port mirroring with next-hop groups. All traffic enters the monitoring router at interface **ge-1/0/0**. A firewall filter counts and port-mirrors all incoming packets to a Tunnel Services PIC. A second filter is applied to the tunnel interface and splits the traffic into three categories: HTTP traffic, FTP traffic, and all other traffic. The three types of traffic are assigned to three separate next-hop groups. Each next-hop group contains a unique pair of exit interfaces that lead to different groups of packet analyzers and flow servers.





**NOTE:** Instances enabled to mirror packets to different destinations from the same PFE, also use different sampling parameters for each instance. When we configure Layer2 Port-mirroring with both global port-mirroring and instance based port-mirroring, PIC level instances will override FPC level and the FPC level will override the Global instance.

```
[edit]
interfaces {
 ge-1/0/0 { # This is the input interface where packets enter the router.
 unit 0 {
 family inet {
 filter {
 input mirror_pkts; # Here is where you apply the first filter.
 }
 address 10.11.1.1/24;
 }
 }
 }
 ge-1/1/0 { # This is an exit interface for HTTP packets.
 unit 0 {
 family inet {
 address 10.12.1.1/24;
 }
 }
 }
 ge-1/2/0 { # This is an exit interface for HTTP packets.
 unit 0 {
 family inet {
 address 10.13.1.1/24;
 }
 }
 }
 so-0/3/0 { # This is an exit interface for FTP packets.
 unit 0 {
 family inet {
 address 10.1.1.1/30;
 }
 }
 }
 so-4/3/0 { # This is an exit interface for FTP packets.
 unit 0 {
 family inet {
 address 10.2.2.1/30;
 }
 }
 }
 so-7/0/0 { # This is an exit interface for all remaining packets.
 unit 0 {
 family inet {
 address 10.5.5.1/30;
 }
 }
 }
}
```

```
so-7/0/1 { # This is an exit interface for all remaining packets.
 unit 0 {
 family inet {
 address 10.6.6.1/30;
 }
 }
}
vt-3/3/0 { # The tunnel interface is where you send the port-mirrored traffic.
 unit 0 {
 family inet;
 }
 unit 1 {
 family inet {
 filter {
 input collect_pkts; # This is where you apply the second firewall filter.
 }
 }
 }
}
forwarding-options {
 port-mirroring { # This is required when you configure next-hop groups.
 family inet {
 input {
 rate 1; # This port-mirrors all packets (one copy for every packet received).
 }
 output { # Sends traffic to a tunnel interface to enable multipoint mirroring.
 interface vt-3/3/0.1;
 no-filter-check;
 }
 }
 }
 next-hop-group ftp-traffic { # Point-to-point interfaces require you to specify the
 interface so-4/3/0.0; # interface name.
 interface so-0/3/0.0;
 }
 next-hop-group http-traffic { # Configure a next hop for all multipoint interfaces.
 interface ge-1/1/0.0 {
 next-hop 10.12.1.2;
 }
 interface ge-1/2/0.0 {
 next-hop 10.13.1.2;
 }
 }
 next-hop-group default-collect {
 interface so-7/0/0.0;
 interface so-7/0/1.0;
 }
}
firewall {
 family inet {
 filter mirror_pkts { # Apply this filter to the input interface.
 term catch_all {
 then {
 count input_mirror_pkts;
 port-mirror; # This action sends traffic to be copied and port-mirrored.
 }
 }
 }
 }
}
```

```

 }
 }
}
filter collect_pkts { # Apply this filter to the tunnel interface.
 term ftp-term { # This term sends FTP traffic to an FTP next-hop group.
 from {
 protocol ftp;
 }
 then next-hop-group ftp-traffic;
 }
 term http-term { # This term sends HTTP traffic to an HTTP next-hop group.
 from {
 protocol http;
 }
 then next-hop-group http-traffic;
 }
 term default { # This sends all remaining traffic to a final next-hop group.
 then next-hop-group default-collectors;
 }
}
}
}

```

- Related Documentation**
- [Understanding Port Mirroring on page 931](#)
  - [Configuring Port Mirroring on page 931](#)



## PART 17

# Real-Time Performance Monitoring and Video Monitoring Services

- [Monitoring Traffic Using Real-Time Performance Monitoring on page 957](#)
- [Testing the Performance of Network Devices Using RFC 2544-Based Benchmarking on page 983](#)
- [Tracking Streaming Media Traffic Using Inline Video Monitoring on page 1043](#)



# Monitoring Traffic Using Real-Time Performance Monitoring

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [Configuring RPM Probes on page 959](#)
- [Configuring RPM Receiver Servers on page 963](#)
- [Limiting the Number of Concurrent RPM Probes on page 964](#)
- [Configuring RPM Timestamping on page 964](#)
- [Configuring TWAMP on page 968](#)
- [Configuring BGP Neighbor Discovery Through RPM on page 971](#)
- [Examples: Configuring BGP Neighbor Discovery Through RPM on page 973](#)
- [Tracing RPM Operations on page 975](#)
- [Examples: Configuring Real-Time Performance Monitoring on page 977](#)
- [Enabling RPM for the Junos OS extension-provider package on page 981](#)

## Real-Time Performance Monitoring Services Overview

---

Real-Time Performance Monitoring (RPM) enables you to configure active probes to track and monitor traffic. Probes collect packets per destination and per application, including PING Internet Control Message Protocol (ICMP) packets, User Datagram Protocol and Transmission Control Protocol (UDP/TCP) packets with user-configured ports, user-configured Differentiated Services code point (DSCP) type-of-service (ToS) packets, and Hypertext Transfer Protocol (HTTP) packets. RPM provides Management Information Base (MIB) support with extensions for RFC 2925, *Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations*.

You can also configure RPM services to determine automatically whether a path exists between a host router and its configured BGP neighbors. You can view the results of the discovery using an SNMP client. Results are stored in **pingResultsTable**, **jnxPingResultsTable**, **jnxPingProbeHistoryTable**, and **pingProbeHistoryTable**.

Probe configuration and probe results are supported by the command-line interface (CLI) and SNMP.

The following probe types are supported with DSCP marking:

- ICMP echo
- ICMP timestamp
- HTTP get (not available for BGP RPM services)
- UDP echo
- TCP connection
- UDP timestamp

With probes, you can monitor the following:

- Minimum round-trip time
- Maximum round-trip time
- Average round-trip time
- Standard deviation of the round-trip time
- Jitter of the round-trip time—The difference between the minimum and maximum round-trip time

One-way measurements for ICMP timestamp probes include the following:

- Minimum, maximum, standard deviation, and jitter measurements for egress and ingress times
- Number of probes sent
- Number of probe responses received
- Percentage of lost probes



**NOTE:** Timestamping is not supported on PTX Series Packet Transport Routers.

---

You can configure the following RPM thresholds:

- Round-trip time
- Ingress/egress delay
- Standard deviation
- Jitter
- Successive lost probes
- Total lost probes (per test)

Support is also implemented for user-configured CoS classifiers and for prioritization of RPM packets over regular data packets received on an input interface.



- Related Documentation**
- [Configuring BGP Neighbor Discovery Through RPM on page 971](#)
  - [\[edit services rpm\] Hierarchy Level on page 1631](#)
  - [Examples: Configuring BGP Neighbor Discovery Through RPM on page 973](#)

## Configuring RPM Probes

The owner name and test name identifiers of an RPM probe together represent a single RPM configuration instance. When you specify the test name, you also can configure the test parameters.

To configure the probe owner, test name, and test parameters, include the **probe** statement at the **[edit services rpm]** hierarchy level:

```
probe owner {
 test test-name {
 data-fill data;
 data-size size;
 destination-interface interface-name;
 destination-port port;
 dscp-code-point dscp-bits;
 hardware-timestamp;
 history-size size;
 moving-average-size number;
 one-way-hardware-timestamp;
 probe-count count;
 probe-interval seconds;
 probe-type type;
 routing-instance instance-name;
 source-address address;
 target (url url | address address);
 test-interval interval;
 thresholds thresholds;
 traps traps;
 }
}
```

Keep the following points in mind when you configure RPM clients and RPM servers:

- You cannot configure an RPM client that is PIC-based and an RPM server that is based on either the Packet Forwarding Engine or Routing Engine to receive the RPM probes.
- You cannot configure an RPM client that is Packet Forwarding Engine-based and an RPM server that receives the RPM probes to be on the PIC or Routing Engine.
- The RPM client and RPM server must be located on the same type of module. For example, if the RPM client is PIC-based, the RPM server must also be PIC-based, and if the RPM server is Packet Forwarding Engine-based, the RPM client must also be Packet Forwarding Engine-based.

- To specify a probe owner, include the **probe** statement at the **[edit services rpm]** hierarchy level. The probe owner identifier can be up to 32 characters in length.
- To specify a test name, include the **test** statement at the **[edit services rpm probe owner]** hierarchy level. The test name identifier can be up to 32 characters in length. A test represents the range of probes over which the standard deviation, average, and jitter are calculated.
- To specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes, include the **data-fill** statement at the **[edit services rpm probe owner]** hierarchy level. The value can be a hexadecimal value. The **data-fill** statement is not valid with the **http-get** or **http-metadata-get** probe types.
- To specify the size of the data portion of ICMP probes, include the **data-size** statement at the **[edit services rpm probe owner]** hierarchy level. The size can be from 0 through 65400 and the default size is 0. The **data-size** statement is not valid with the **http-get** or **http-metadata-get** probe types.



**NOTE:** If you configure the hardware timestamp feature (see [“Configuring RPM Timestamping” on page 964](#)):

- The **data-size** default value is 32 bytes and 32 is the minimum value for explicit configuration. The UDP timestamp probe type is an exception; it requires a minimum data size of 44 bytes.
  - The **data-size** must be at least 100 bytes smaller than the default MTU of the interface of the RPM client interface.
- 
- On M Series and T Series routers, you configure the **destination-interface** statement to enable hardware timestamping of RPM probe packets. You specify an **sp-** interface to have the AS or Multiservices PIC add the hardware timestamps; for more information, see [“Configuring RPM Timestamping” on page 964](#). You can also include the **one-way-hardware-timestamp** statement to enable one-way delay and jitter measurements.
  - To specify the User Datagram Protocol (UDP) port or Transmission Control Protocol (TCP) port to which the probe is sent, include the **destination-port** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. The **destination-port** statement is used only for the UDP and TCP probe types. The value can be 7 or from 49160 through 65535.

When you configure either **probe-type udp-ping** or **probe-type udp-ping-timestamp** along with hardware timestamping, the value for the **destination-port** can be only 7. A constraint check prevents you from configuring any other value for the destination port in this case. This constraint does not apply when you are using one-way hardware timestamping.

- To specify the value of the Differentiated Services (DiffServ) field within the IP header, include the **dscp-code-point** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. The DiffServ code point (DSCP) bits value can be set to a valid 6-bit pattern; for example, 001111. It also can be set using an alias configured at

the **[edit class-of-service code-point-aliases dscp]** hierarchy level. The default is 000000.

- To specify the number of stored history entries, include the **history-size** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. Specify a value from 0 to 512. The default is 50.
- To specify a number of samples for making statistical calculations, include the **moving-average-size** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. Specify a value from 0 through 255.
- To specify the number of probes within a test, include the **probe-count** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. Specify a value from 1 through 15.
- To specify the time to wait between sending packets, include the **probe-interval** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. Specify a value from 1 through 255 seconds.
- To specify the packet and protocol contents of the probe, include the **probe-type** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. The following probe types are supported:
  - **http-get**—Sends a Hypertext Transfer Protocol (HTTP) get request to a target URL.
  - **http-metadata-get**—Sends an HTTP get request for metadata to a target URL.
  - **icmp-ping**—Sends ICMP echo requests to a target address.
  - **icmp-ping-timestamp**—Sends ICMP timestamp requests to a target address.
  - **tcp-ping**—Sends TCP packets to a target.
  - **udp-ping**—Sends UDP packets to a target.
  - **udp-ping-timestamp**—Sends UDP timestamp requests to a target address.

The following probe types support hardware timestamping of probe packets: **icmp-ping**, **icmp-ping-timestamp**, **udp-ping**, **udp-ping-timestamp**.



**NOTE:** Some probe types require additional parameters to be configured. For example, when you specify the **tcp-ping** or **udp-ping** option, you must configure the destination port using the **destination-port** statement. The **udp-ping-timestamp** option requires a minimum data size of 12; any smaller data size results in a commit error. The minimum data size for TCP probe packets is 1.

When you configure either **probe-type udp-ping** or **probe-type udp-ping-timestamp** along with the **one-way-hardware-timestamp** command, the value for the **destination-port** can be only 7. A constraint check prevents you for configuring any other value for the destination port in this case.

- To specify the routing instance used by ICMP probes, include the **routing-instance** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. The default routing instance is Internet routing table **inet.0**.
- To specify the source IP address used for ICMP probes, include the **source-address** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. If the source IP address is not one of the router's assigned addresses, the packet will use the outgoing interface's address as its source.
- To specify the destination address used for the probes, include the **target** statement at the **[edit services rpm probe owner test test-name]** hierarchy level.
  - For HTTP probe types, specify a fully formed URL that includes **http://** in the URL address.
  - For all other probe types, specify an IP version 4 (IPv4) address for the target host.
- To specify the time to wait between tests, include the **test-interval** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. Specify a value from 1 through 86400 seconds.



**NOTE:** Starting with Junos OS Release 15.1, the minimum period for which the RPM client waits between two tests is modified to be 1 second instead of 0 seconds. Also, if you do not configure the test interval, the default value is 1 second.

---

- To specify thresholds used for the probes, include the **thresholds** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. A system log message is generated when the configured threshold is exceeded. Likewise, an SNMP trap (if configured) is generated when a threshold is exceeded. The following options are supported:
  - **egress-time**—Measures maximum source-to-destination time per probe.
  - **ingress-time**—Measures maximum destination-to-source time per probe.
  - **jitter-egress**—Measures maximum source-to-destination jitter per test.
  - **jitter-ingress**—Measures maximum destination-to-source jitter per test.
  - **jitter-rtt**—Measures maximum jitter per test, from 0 through 60000000 microseconds.
  - **rtt**—Measures maximum round-trip time per probe, in microseconds.
  - **std-dev-egress**—Measures maximum source-to-destination standard deviation per test.
  - **std-dev-ingress**—Measures maximum destination-to-source standard deviation per test.
  - **std-dev-rtt**—Measures maximum standard deviation per test, in microseconds.

- **successive-loss**—Measures successive probe loss count, indicating probe failure.
- **total-loss**—Measures total probe loss count indicating test failure, from 0 through 15.
- Traps are sent if the configured threshold is met or exceeded. To set the trap bit to generate traps, include the **traps** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. The following options are supported:
  - **egress-jitter-exceeded**—Generates traps when the jitter in egress time threshold is met or exceeded.
  - **egress-std-dev-exceeded**—Generates traps when the egress time standard deviation threshold is met or exceeded.
  - **egress-time-exceeded**—Generates traps when the maximum egress time threshold is met or exceeded.
  - **ingress-jitter-exceeded**—Generates traps when the jitter in ingress time threshold is met or exceeded.
  - **ingress-std-dev-exceeded**—Generates traps when the ingress time standard deviation threshold is met or exceeded.
  - **ingress-time-exceeded**—Generates traps when the maximum ingress time threshold is met or exceeded.
  - **jitter-exceeded**—Generates traps when the jitter in round-trip time threshold is met or exceeded.
  - **probe-failure**—Generates traps for successive probe loss thresholds crossed.
  - **rtt-exceeded**—Generates traps when the maximum round-trip time threshold is met or exceeded.
  - **std-dev-exceeded**—Generates traps when the round-trip time standard deviation threshold is met or exceeded.
  - **test-completion**—Generates traps when a test is completed.
  - **test-failure**—Generates traps when the total probe loss threshold is met or exceeded.

**Related  
Documentation**

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [Examples: Configuring Real-Time Performance Monitoring on page 977](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)

---

## Configuring RPM Receiver Servers

The RPM TCP and UDP probes are proprietary to Juniper Networks and require a receiver to receive the probes. To configure a server to receive the probes, include the **probe-server** statement at the **[edit services rpm]** hierarchy level:

```
[edit services rpm]
probe-server {
```

```
tcp {
 destination-interface interface-name;
 port number;
}
udp {
 port number;
}
}
```

The port number specified for the UDP and TCP server can be 7 or from 49160 through 65535.



**NOTE:** The `destination-interface` statement is not supported on PTX Series Packet Transport Routers.

When you configure either `probe-type udp-ping` or `probe-type udp-ping-timestamp` along with the `one-way-hardware-timestamp` command, the value for the `destination-port` can be only 7. A constraint check prevents you for configuring any other value for the destination port in this case.

**Related  
Documentation**

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)
- [Examples: Configuring Real-Time Performance Monitoring on page 977](#)

---

## Limiting the Number of Concurrent RPM Probes

To configure the maximum number of concurrent probes allowed, include the `probe-limit` statement at the `[edit services rpm]` hierarchy level:

```
probe-limit limit;
```

Specify a limit from 1 through 500. The default maximum number is 100.

**Related  
Documentation**

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)
- [Examples: Configuring Real-Time Performance Monitoring on page 977](#)

---

## Configuring RPM Timestamping

To account for latency in the communication of probe messages, you can enable timestamping of the probe packets. You can timestamp the following RPM probe types: `icmp-ping`, `icmp-ping-timestamp`, `udp-ping`, and `udp-ping-timestamp`.

On M Series and T Series routers with an Adaptive Services (AS) or Multiservices PIC, and on EX Series switches with a Multiservices DPC, and on EX Series switches, you can enable hardware timestamping of RPM probe messages. The timestamp is applied on both the RPM client router (the router or switch that originates the RPM probes) and the RPM probe server and applies only to IPv4 traffic. It is supported on the following:

- Layer 2 services package on all Multiservices PICs and DPCs.
- Layer 3 service package on AS and Multiservices PICs and Multiservices DPCs.
- Extension-provider services package on M Series, MX Series, and T Series services PICs that support the Extension-Provider packages (In Junos OS releases earlier than 12.3, the extension-provider packages were variously referred to as Junos Services Framework (JSF), MP-SDK, and eJunos.)
- Layer 2, Layer 3, SDK Services, and PFE RPM timestamping interoperate with each other. Here, the RPM client can be on the Layer 3 **sp-** interface and the RPM server can be on an SDK Services package.



**NOTE:** Hardware timestamping is not supported on PTX Series Packet Transport Routers.

Two-way timestamping is available on **sp-** and **ms-** interfaces. To configure two-way timestamping on M Series and T Series routers, include the **destination-interface** statement at the **[edit services rpm probe probe-owner test test-name]** hierarchy level:

```
destination-interface sp-fpc/pic/port.logical-unit
destination-interface ms-fpc/pic/port.logical-unit
```

Specify the RPM client router and the RPM server router on the adaptive services logical interface or the multiservices interface by including the **rpm** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level:

```
rpm (client | server);
```

The logical interface must be dedicated to the RPM task. It requires configuration of the **family inet** statement and a **/32** address, as shown in the example. This configuration is also needed for other services such as NAT and stateful firewall. You cannot configure RPM service on **unit 0** because RPM requires a dedicated logical interface; the same unit cannot support both RPM and other services. Because active flow monitoring requires **unit 0**, but RPM can function on any logical interface, a constraint check prevents you from committing an RPM configuration there.



**NOTE:** If you configure RPM timestamping on an AS PIC, you cannot configure the **source-address** statement at the **[edit services rpm probe probe-name test test-name]** hierarchy level.

On MX Series routers, on M-320 routers using the Enhanced Queuing MPC, and on EX Series switches, you include the **hardware-timestamp** statement at the **[edit services rpm**

**probe *probe-name* test *test-name*]** hierarchy level to specify that the probes are to be timestamped in the Packet Forwarding Engine host processor:

**hardware-timestamp;**

On the client side, these probes are timestamped in the Packet Forwarding Engine host processor on the egress DPC on the MX or M-320 Series router or EX Series switch originating the RPM probes (RPM client). On the responder side (RPM server), the RPM probes to be timestamped are handled by the Packet Forwarding Engine host processor, which generates the response instead of the RPM process. The RPM probes are timestamped only on the router that originates them (RPM client). As a result, only round-trip time is measured for these probes.

When using the **hardware-timestamp**, the **data-size** value for the probe must be at least 100 bytes smaller than the default MTU of the interface of the RPM client interface (see [“Configuring RPM Probes” on page 959](#)). If hardware timestamping of RPM probe messages is enabled, the maximum data size that you can configure by using the data-size statement is limited to 1400.



**NOTE:** The Packet Forwarding Engine-based RPM feature does not support any stateful firewall configurations. If you need to combine RPM timestamping with a stateful firewall, you should use the interface-based RPM timestamping service described earlier in this section. Multiservices DPCs support stateful firewall processing as well as RPM timestamping.

---

To configure one-way timestamping, you must also include the **one-way-hardware-timestamp** statement at the **[edit services rpm probe *probe-owner* test *test-name*]** hierarchy level:

**one-way-hardware-timestamp;**





**NOTE:** If you configure RPM probes for a services interface (sp-), you need to announce local routes in a specific way for the following routing protocols:

- For OSPF, you can announce the local route by including the services interface in the OSPF area. To configure this setting, include the interface `sp-fpc/pic/port` statement at the [edit protocols ospf area *area-number*] hierarchy level.
- For BGP and IS-IS, you must export interface routes and create a policy that accepts the services interface local route. To export interface routes, include the point-to-point and lan statements at the [edit routing-options interface-routes family inet export] hierarchy level. To configure an export policy that accepts the services interface local route, include the protocol local, rib inet.0, and route-filter `sp-interface-ip-address/32` exact statements at the [edit policy-options policy-statement *policy-name* term *term-name* from] hierarchy level and the accept action at the [edit policy-options policy-statement *policy-name* term *term-name* then] hierarchy level. For the export policy to take effect, apply the policy to BGP or IS-IS with the export *policy-name* statement at the [edit protocols *protocol-name*] hierarchy level.

For more information about these configurations, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices* or the *Junos OS Routing Protocols Library for Routing Devices*.

Routing the probe packets through the adaptive services or Multiservices PIC also enables you to filter the probe packets to particular queues. The following example shows the RPM configuration and the filter that specifies queuing:

```
services rpm {
 probe p1 {
 test t1 {
 probe-type icmp-ping;
 target address 10.8.4.1;
 probe-count 10;
 probe-interval 10;
 test-interval 10;
 dscp-code-points af11;
 data-size 100;
 destination-interface sp-1/2/0.0;
 }
 }
}
firewall {
 filter f1 {
 term t1 {
 from {
 dscp af11;
 }
 then {
 forwarding-class assured-forwarding;
 }
 }
 }
}
```

```
 }
 }
}
interfaces sp-1/2/0 {
 unit 2 {
 rpm client;
 family inet {
 address 10.8.4.2/32;
 filter {
 input f1;
 }
 }
 }
}
interfaces sp-1/2/1 {
 unit 2 {
 rpm server;
 family inet {
 address 10.8.3.2/32;
 filter {
 input f1;
 }
 }
 }
}
```

For more information about firewall filters, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*; for more information about queuing, see the *Class of Service Feature Guide for Routing Devices*.

#### Related Documentation

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)
- [Examples: Configuring Real-Time Performance Monitoring on page 977](#)

---

## Configuring TWAMP

You can configure the Two-Way Active Measurement Protocol (TWAMP) on all M Series and T Series routers that support Multiservices PICs (running in either Layer 2 or Layer 3 mode), and on MX Series routers. Only the responder (server) side of TWAMP is supported.



**NOTE:** TWAMP is not supported on EX Series switches and PTX Series Packet Transport Routers.

For more information on TWAMP, see RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP)*.

To configure TWAMP properties, include the **twamp** statement at the [\[edit services rpm\]](#) hierarchy level:

```
[edit services rpm]
```

```

twamp {
 server {
 client-list list-name {
 [address address];
 }
 authentication-mode mode;
 max-connection-duration hours;
 maximum-connections count;
 maximum-connections-per-client count;
 maximum-sessions count;
 maximum-sessions-per-connection count;
 port number;
 routing-instance-list {
 instance-name {
 port number;
 }
 }
 server-inactivity-timeout minutes;
 }
}

```

The TWAMP configuration process includes the following tasks:

- [Configuring TWAMP Interfaces on page 969](#)
- [Configuring TWAMP Servers on page 969](#)

## Configuring TWAMP Interfaces

To specify the service PIC logical interface that provides the TWAMP service, include the **twamp-server** statement at the **[edit interfaces sp-fpc/pic/port unit logical-unit-number]** hierarchy level:

```
twamp-server;
```



**NOTE:** On MX Series routers that do not include a Multiservices DPC, you can configure the **twamp-server** statement on any interface (for example, **ge-1/0/1.10**). It is not necessary to configure this statement on a service interface (**sp-** or **ms-**) but you do need to include it in the configuration to activate the TWAMP reflector functionality.

## Configuring TWAMP Servers

You can specify a number of TWAMP server properties, some of which are optional, by including the **server** statement at the **[edit services rpm twamp]** hierarchy level:

```

[edit services rpm twamp]
server {
 client-list list-name {
 [address address];
 }
 authentication-mode mode;
 max-connection-duration hours;
 maximum-connections count;

```

```
maximum-connections-per-client count;
maximum-sessions count;
maximum-sessions-per-connection count;
port number;
routing-instance-list {
 instance-name {
 port number;
 }
}
server-inactivity-timeout minutes;
}
```

The preceding configuration settings that are described define a TWAMP server on the router that enables a TWAMP client to connect to the server using any media interface IP address such as a **ge-** interface. In such a scenario, the router functions as a TWAMP server and timestamping is performed in the ukernel of the media-facing FPC.

To configure an inline TWAMP server, which causes timestamping to be performed as part of the inline services (**si-**) interface processing, configure the amount of bandwidth reserved on each Packet Forwarding Engine for tunnel traffic using inline services by including the **bandwidth (1g | 10g)** statement at the **[edit chassis fpc slot-number pic number inline-services]** hierarchy level and specify the service PIC logical interface that provides the TWAMP service by including the **twamp-server** statement at the **[edit interfaces sp-fpc/pic/port unit logical-unit- number family inet]** hierarchy level.

- To specify the list of allowed control client hosts that can connect to this server, include the **client-list** statement at the **[edit services rpm twamp server]** hierarchy level. Each value you include must be a Classless Interdomain Routing (CIDR) address (IP address plus mask) that represents a network of allowed hosts. You can include multiple client lists, each of which can contain a maximum of 64 entries. You must configure at least one client address to enable TWAMP.
- You must specify the authentication mode by including the **authentication-mode** statement at the **[edit services rpm twamp server]** hierarchy level. There is no default value. You can configure **authenticated** or **encrypted** mode, based on RFC 4656; if there is no authentication or encryptions mode specified, you should set the value to **none**. This statement is required in the TWAMP configuration.
- To specify the inactivity timeout period in seconds, include the **inactivity-timeout** statement at the **[edit services rpm twamp server]** hierarchy level. By default, the value is **1800**; the range is 0 through 3600 seconds.
- To specify the maximum number of concurrent connections the server can have to client hosts, include the **maximum-connections** statement at the **[edit services rpm twamp server]** hierarchy level. The allowed range of values is 1 through 1000 and the default value is 64. You can also limit the number of connections the server can make to a particular client host by including the **maximum-connections-per-client** statement. The allowed range of values is 1 through 500 and the default value is 64.
- To specify the maximum number of sessions the server can have running at one time, include the **maximum-sessions** statement at the **[edit services rpm twamp server]** hierarchy level. The allowed range of values is 1 through 2048 and the default value is

64. You can also limit the number of sessions the server can have on a single connection by including the **maximum-sessions-per-connection** statement.

- To specify the TWAMP server listening port, include the **port** statement at the **[edit services rpm twamp server]** hierarchy level. The range is 1 through 65,535.
- To specify the server inactivity timeout period in minutes, include the **server-inactivity-timeout** statement at the **[edit services rpm twamp server]** hierarchy level. The range is 0 through 30 minutes.
- To specify the TWAMP servers on specific routing instances, instead of associating the TWAMP server at the system-level to apply to all routing instances configured on a router, include the **routing-instance-list instance-name port port-number** statement at the **[edit services rpm twamp server]** hierarchy level. The port number of the specified routing instance is used for TWAMP probes that are received by a TWAMP server. The default routing instance is Internet routing table inet.0. If you do not specify a routing instance, the TWAMP probe applies to all routing instances. To apply the TWAMP probe to only the default routing instance, you must explicitly set the value of instance-name to default. If an interface is not part of any routing instance, the default port is used for TWAMP probes. You can configure up to 100 routing instances for a TWAMP server.

## Configuring BGP Neighbor Discovery Through RPM

BGP neighbors can be configured at the following hierarchy levels:

- **[edit protocols bgp group group-name]**—Default logical system and default routing instance.
- **[edit routing-instances instance-name protocols bgp group group-name]**—Default logical system with a specified routing instance.
- **[edit logical-systems logical-system-name protocols bgp group group-name]**—Configured logical system and default routing instance.
- **[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group group-name]**—Configured logical system with a specified routing instance.

When you configure BGP neighbor discovery through RPM, if you do not specify a logical system, the RPM probe applies to configured BGP neighbors for all logical systems. If you do not specify a routing instance, the RPM probe applies to configured BGP neighbors in all routing instances. You can explicitly configure RPM probes to apply only to the default logical system, the default routing instance, or to a particular logical system or routing instance.

To configure BGP neighbor discovery through RPM, configure the probe properties at the **[edit services rpm bgp]** hierarchy:

```
data-fill data;
data-size size;
destination-port port;
history-size size;
logical-system logical-system-name [routing-instances routing-instance-name];
```

**moving-average-size** *number*;  
**probe-count** *count*;  
**probe-interval** *seconds*;  
**probe-type** *type*;  
**routing-instances** *instance-name*;  
**test-interval** *interval*;

- To specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes, include the **data-fill** statement at the **[edit services rpm bgp]** hierarchy level. The value can be a hexadecimal value.
- To specify the size of the data portion of ICMP probes, include the **data-size** statement at the **[edit services rpm bgp]** hierarchy level. The size can be from 0 through 65400 and the default size is 0.
- To specify the User Datagram Protocol (UDP) port or Transmission Control Protocol (TCP) port to which the probe is sent, include the **destination-port** statement at the **[edit services rpm bgp]** hierarchy level. The **destination-port** statement is used only for the UDP and TCP probe types. The value can be 7 or from 49160 through 65535.
- To specify the number of stored history entries, include the **history-size** statement at the **[edit services rpm bgp]** hierarchy level. Specify a value from 0 to 512. The default is 50.
- To specify the logical system used by ICMP probes, include the **logical-system** *logical-system-name* statement at the **[edit services rpm bgp]** hierarchy level. If you do not specify a logical system, the RPM probe applies to configured BGP neighbors for all logical systems. To apply the probe to only the default logical system, you must set the value of *logical-system-name* to null.
- To specify a number of samples for making statistical calculations, include the **moving-average-size** statement at the **[edit services rpm bgp]** hierarchy level. Specify a value from 0 through 255.
- To specify the number of probes within a test, include the **probe-count** statement at the **[edit services rpm bgp]** hierarchy level. Specify a value from 1 through 15.
- To specify the time to wait between sending packets, include the **probe-interval** statement at the **[edit services rpm bgp]** hierarchy level. Specify a value from 1 through 255 seconds.
- To specify the packet and protocol contents of the probe, include the **probe-type** statement at the **[edit services rpm bgp]** hierarchy level. The following probe types are supported:
  - **icmp-ping**—Sends ICMP echo requests to a target address.
  - **icmp-ping-timestamp**—Sends ICMP timestamp requests to a target address.
  - **tcp-ping**—Sends TCP packets to a target.
  - **udp-ping**—Sends UDP packets to a target.
  - **udp-ping-timestamp**—Sends UDP timestamp requests to a target address.



**NOTE:** Some probe types require additional parameters to be configured. For example, when you specify the `tcp-ping` or `udp-ping` option, you must configure the destination port using the `destination-port port` statement. The `udp-ping-timestamp` option requires a minimum data size of 12; any smaller data size results in a commit error. The minimum data size for TCP probe packets is 1.

- To specify the routing instance used by ICMP probes, include the **routing-instances** statement at the `[edit services rpm bgp]` hierarchy level. The default routing instance is Internet routing table `inet.0`. If you do not specify a routing instance, the RPM probe applies to configured BGP neighbors in all routing instances. To apply the RPM probe to only the default routing instance, you must explicitly set the value of *instance-name* to **default**.
- To specify the time to wait between tests, include the **test-interval** statement at the `[edit services bgp probe]` hierarchy level. Specify a value from 1 through 86400 seconds.



**NOTE:** Starting with Junos OS Release 15.1, the minimum period for which the RPM client waits between two tests is modified to be 1 second instead of 0 seconds. Also, if you do not configure the test interval, the default value is 1 second.

#### Related Documentation

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)
- [Examples: Configuring BGP Neighbor Discovery Through RPM on page 973](#)

## Examples: Configuring BGP Neighbor Discovery Through RPM

Configure BGP neighbor discovery through RPM for all logical systems and all routing instances:

```
[edit services rpm]
bgp {
 probe-type icmp-ping;
 probe-count 5;
 probe-interval 1;
 test-interval 60;
 history-size 10;
 data-size 255;
 data-fill 0123456789;
}
```

Configure BGP neighbor discovery through RPM for only the following logical systems and routing instances: **LS1/RI1**, **LS1/RI2**, **LS2**, and **RI3**:

```
[edit services rpm]
bgp {
```

```
probe-type icmp-ping;
probe-count 5;
probe-interval 1;
test-interval 60;
history-size 10;
data-size 255;
data-fill 0123456789;
logical-system {
 LS1 {
 routing-instances {
 RI1;
 RI2;
 }
 }
 LS2;
}
routing-instance {
 RI3;
}
}
```



**NOTE:** The `logical-system` statement is not supported on PTX Series Packet Transport Routers.

Configure BGP neighbor discovery through RPM for only the default logical system and default routing instance:

```
[edit services rpm]
bgp {
 probe-type icmp-ping;
 probe-count 5;
 probe-interval 1;
 test-interval 60;
 history-size 10;
 data-size 255;
 data-fill 0123456789;
 logical-system {
 null {
 routing-instances {
 default;
 }
 }
 }
}
```

**Related  
Documentation**

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [Configuring BGP Neighbor Discovery Through RPM on page 971](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)



## Tracing RPM Operations

Tracing operations track all RPM operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

By default, no events are traced. If you include the **traceoptions** statement at the **[edit services rpm]** hierarchy level, the default tracing behavior is the following:

- Important events are logged in a file called **rmopd** located in the **/var/log** directory.
- When the log file reaches 128 kilobytes (KB), it is renamed **rmopd.0**, then **rmopd.1**, and so on, until there are three trace files. Then the oldest trace file (**rmopd.2**) is overwritten. (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)
- Log files can be accessed only by the user who configures the tracing operation.

You can change this default behavior by using the **traceoptions** statements. Changing the defaults is described in the following sections:

1. [Configuring the RPM Log File Name on page 975](#)
2. [Configuring the Number and Size of RPM Log Files on page 975](#)
3. [Configuring Access to the Log File on page 976](#)
4. [Configuring a Regular Expression for Lines to Be Logged on page 976](#)
5. [Configuring the Trace Operations on page 976](#)

### Configuring the RPM Log File Name

By default, the name of the file that records RPM trace output is **rmopd**. To specify a different file name:

```
[edit services rpm traceoptions]
user @host set file filename
```

### Configuring the Number and Size of RPM Log Files

To configure the limits on the number and size of RPM trace files:

```
[edit services rpm traceoptions]
user@host set file filename files number size size
```

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

For example, set the maximum file size to 2 MB, and the maximum number of files to 20 for a log file named **rpmtrace**:

```
[edit services rpm traceoptions]
user@host set file rpmtrace files 20 size 2MB
```

When the **rpmtrace** file reaches 2 MB, it is renamed **rpmtrace.0**, and a new file called **rpmtrace** is created. When the new **rpmtrace** reaches 2 MB, **rpmtrace.0** is renamed

**rpmtrace.1** and **rpmtrace** is renamed **rpmtrace.0**. This process repeats until there are 20 trace files. Then the oldest file (**rpmtrace.19**) is overwritten by **rpmtrace.18**.

## Configuring Access to the Log File

By default, log files can be accessed only by the user who configures the tracing operation.

To specify that any user can read all log files:

```
[edit services rpm traceoptions]
user@host set file filename world-readable
```

To explicitly set the default behavior:

```
[edit services rpm traceoptions]
user@host set file filename no-world-readable
```

## Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

To refine the output by specifying a regular expression (regex) to be matched:

```
[edit services rpm traceoptions]
user@host set file filename match regular-expression
```

## Configuring the Trace Operations

By default, if the **traceoptions** configuration is present, only important events are logged. You can configure the trace operations to be logged by including the following statements at the **[edit services rpm traceoptions]** hierarchy level:

```
flag {
 all;
 configuration;
 error;
 ipc;
 ppm;
 statistics
}
```

[Table 7 on page 51](#) describes the meaning of the RPM tracing flags.

**Table 36: RPM Tracing Flags**

Flag	Description	Default Setting
<b>all</b>	Trace all operations.	Off
<b>configuration</b>	Trace configuration events.	Off
<b>error</b>	Trace events related to catastrophic errors in daemon.	Off
<b>ipc</b>	Trace IPC events.	Off

Table 36: RPM Tracing Flags (*continued*)

Flag	Description	Default Setting
<b>ppm</b>	Trace ppm events.	Off
<b>statistics</b>	Trace statistics.	Off

## Examples: Configuring Real-Time Performance Monitoring

Configure an RPM instance identified by the probe name **probe1** and the test name **test1**:

```
[edit services rpm]
probe probe1 {
 test test1 {
 dscp-code-points 001111;
 probe-interval 1;
 probe-type icmp-ping;
 target address 172.17.20.182;
 test-interval 20;
 thresholds rtt 10;
 traps rtt-exceeded;
 }
}
probe-server {
 tcp {
 destination-interface lt-0/0/0.0
 port 50000;
 }
 udp {
 destination-interface lt-0/0/0.0
 port 50001;
 }
}
probe-limit 200;
```

Configure packet classification, using **lt-** interfaces to send the probe packets to a logical tunnel input interface. By sending the packet to the logical tunnel interface, you can configure regular and multifield classifiers, firewall filters, and header rewriting for the probe packets. To use the existing tunnel framework, the **dlci** and **encapsulation** statements must be configured.

```
[edit services rpm]
probe p1 {
 test t1 {
 probe-type icmp-ping;
 target address 10.8.4.1;
 probe-count 10;
 probe-interval 10;
 test-interval 10;
 source-address 10.8.4.2;
 dscp-code-points ef;
 data-size 100;
 destination-interface lt-0/0/0.0;
 }
}
```

```
}
[edit interfaces]
lt-0/0/0 {
 unit 0 {
 encapsulation frame-relay;
 dlci 10;
 peer-unit 1;
 family inet;
 }
 unit 1 {
 encapsulation frame-relay;
 dlci 10;
 peer-unit 0;
 family inet;
 }
}
[edit class-of-service]
interfaces {
 lt-0/0/0 {
 unit 1 {
 classifiers {
 dscp default;
 }
 }
 }
}
```

Configure an input filter on the interface on which the RPM probes are received. This filter enables prioritization of the received RPM packets, separating them from the regular data packets received on the same interface.

```
[edit firewall]
filter recos {
 term recos {
 from {
 source-address {
 10.8.4.1/32;
 }
 destination-address {
 10.8.4.2/32;
 }
 }
 then {
 loss-priority high;
 forwarding-class network-control;
 }
 }
}
[edit interfaces]
fe-5/0/0 {
 unit 0 {
 family inet {
 filter {
 input recos;
 }
 address 10.8.4.2/24;
 }
 }
}
```

```

 }
 }
}

```

Configure an RPM instance and enable RPM for the extension-provider packages on the adaptive services interface:

```

[edit services rpm]
probe probe1 {
 test test1 {
 data-size 1024;
 data-fill 0;
 destination-interface ms-1/2/0.10;
 dscp-code-points 001111;
 probe-count 10;
 probe-interval 1;
 probe-type icmp-ping;
 target address 172.17.20.182;
 test-interval 20;
 thresholds rtt 10;
 traps rtt-exceeded;
 }
}
[edit interfaces]
ms-1/2/0 {
 unit 0 {
 family inet;
 }
 unit 10 {
 rpm client;
 family inet {
 address 1.1.1.1/32;
 }
 }
}
[edit chassis]
fpc 1 {
 pic 2 {
 adaptive-services {
 service-package {
 extension-provider {
 control-cores 1;
 data-cores 1;
 object-cache-size 512;
 policy-db-size 64;
 package jservices-rpm;
 syslog {
 daemon any;
 }
 }
 }
 }
 }
}
}
}
}

```



**NOTE:** TWAMP is not supported on PTX Series Packet Transport Routers.

Configure the minimum statements necessary to enable TWAMP:

```
[edit services]
rpm {
 twamp {
 server {
 authentication-mode none;
 port 10000; # Twamp server's listening port
 client-list LIST-1 { # LIST-1 is the name of the client-list. Multiple lists can be
 configured.
 address {
 20.0.0.2/30; # IP address of the control client.
 }
 }
 }
 }
}
[edit interfaces sp-5/0/0]
unit 0 {
 family inet;
}
unit 10 {
 rpm {
 twamp-server; # You must configure a separate logical interface on the service PIC
 interface for the TWAMP server.
 }
 family inet {
 address 50.50.50.50/32; # This address must be a host address with a 32-bit mask.
 }
}
[edit chassis]
fpc 5 {
 pic 0 {
 adaptive-services {
 service-package layer-2; # Configure the service PIC to run in Layer 2 mode.
 }
 }
}
```

Configure additional TWAMP settings:

```
[edit services]
rpm {
 twamp {
 server {
 maximum-sessions 5;
 maximum-sessions-per-connection 2;
 maximum-connections 3;
 maximum-connections-per-client 1;
 port 10000;
 server-inactivity-timeout ;
 client-list LIST-1 {
 address {
```

```

 20.0.0.2/30;
 }
}
}
}
}

```

#### Related Documentation

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)
- [Examples: Configuring BGP Neighbor Discovery Through RPM on page 973](#)

## Enabling RPM for the Junos OS extension-provider package

Real-time performance monitoring (RPM), which has been supported on the adaptive services interface, is now supported by the Junos OS extension-provider package. RPM is supported on all platforms and service PICs that support the extension-provider package.



**NOTE:** In Junos OS releases earlier than 12.3, the extension provider package was variously known as MP-SDK, Junos Services Framework (JSF), and eJunos.

To enable RPM for the Junos OS extension-provider package on the adaptive services interface, configure the **object-cache-size**, **policy-db-size**, and **package** statements at the **[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]** hierarchy level. For the extension-provider package, **package-name** in the **package package-name** statement is **jservices-rpm**.

For more information about the extension-provider package, see the *SDK Applications Configuration Guide and Command Reference*.

The following example shows how to enable RPM for the extension-provider package on the adaptive services interface:

```

chassis fpc 1 {
 pic 2 {
 adaptive-services {
 service-package {
 extension-provider {
 control-cores 1;
 data-cores 1;
 object-cache-size 512;
 policy-db-size 64;
 package jservices-rpm;
 syslog daemon any;
 }
 }
 }
 }
}

```

}

**Related  
Documentation**

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)
- [Examples: Configuring Real-Time Performance Monitoring on page 977](#)
- [destination-interface on page 1652](#)



# Testing the Performance of Network Devices Using RFC 2544-Based Benchmarking

- [RFC2544-Based Benchmarking Tests Overview on page 983](#)
- [Layer 2 RFC2544-Based Benchmarking Tests Overview on page 986](#)
- [Supported RFC2544-Based Benchmarking Statements on MX104 Routers on page 988](#)
- [Configuring an RFC 2544-Based Benchmarking Test on page 989](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for Layer 3 IPv4 Services on page 993](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for UNI Direction of Ethernet Pseudowires on page 1001](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for NNI Direction of Ethernet Pseudowires on page 1008](#)
- [Example: Configuring RFC2544-Based Benchmarking Tests for Layer 2 E-LAN Services in Bridge Domains on page 1016](#)

## RFC2544-Based Benchmarking Tests Overview

---

RFC2544 defines a series of tests that can be used to describe the performance characteristics of a network-interconnecting device, such as a router, and outlines specific formats to report the results of the tests. These tests can be used to benchmark interconnected network devices and devise a guideline or a measurement pattern to analyze the health and efficiency of the network devices. These tests are the standard benchmarking tests for Ethernet networks and are known as RFC2544-based benchmarking tests. These tests measure throughput, latency, frame loss rate, and bursty frames. The test methodology enables you to define various parameters such as different frame sizes to be examined (64, 128, 256, 512, 1024, 1280, and 1518 bytes), the test time for each test iteration (10 seconds to 1,728,000 seconds), and the frame format (UDP-over-IP).



**NOTE:** RFC2544-based benchmarking tests support only UDP over IPv4 test traffic.

An RFC2544-based benchmarking test is performed by transmitting test packets from a device that functions as the generator or the initiator (which is also called the originator). These packets are sent to a device that functions as a reflector, which receives and returns the packets to the initiator.

Juniper Networks MX104 3D Universal Edge Routers support only the reflector function and the corresponding benchmarking tests. These tests display only the reflecting benchmarking tests. These benchmarking tests display the results of the test. For instance, in the case of the throughput test, the results display the number of transmitted frames and the number of received frames.

The RFC2544-based benchmarking test methodology assesses different parameters that are defined in service-level agreements (SLAs). By measuring the performance availability, transmission delay, link bursts, and service integrity, a carrier provider can certify that the working parameters of the deployed Ethernet circuit comply with the SLA and other defined policies.

[Table 37 on page 984](#) describes the different network topologies in which the benchmarking test is supported.

**Table 37: Supported Network Topologies for RFC2544 Benchmarking Tests**

Service Type	Traffic Direction	Mode	Initial Release on MX104 Routers	Whether the Benchmarking Test Is Supported
E-Line and E-LAN (family <b>bridge</b> )	(UNI) Egress	Port Port, VLAN	14.2R1 (E-Line and E-LAN family bridge)	Supported
E-Line (family <b>ccc</b> )	Ingress Egress		13.3R1 (E-Line Pseudowire)	Supported
IP Services (family <b>inet</b> )	NNI		13.3R1	Supported



**NOTE:** You can configure a total of four simultaneous active reflection sessions. The four active reflection sessions can be of the same type or can be a combination of the different types of reflection sessions. For instance, you can configure either four IPv4 reflection sessions or two pseudowire reflection sessions, one Layer 2 reflection session, and one IPv4 reflection session. The maximum reflection bandwidth supported is 4Gbps.

[Table 38 on page 985](#) lists the interfaces and the reflection type on which the benchmarking tests are supported.

**Table 38: Supported Interfaces for RFC2544 Benchmarking Tests**

Type of Reflection	Gigabit Interfaces (ge)	Aggregated Interfaces (ae)	10G Interfaces (xe)	Pseudo Interfaces (irb, lt, vt, lo0, and others)
IPv4	Yes	No	No	No
Pseudowire Ingress	Yes	No	No	No
Pseudowire Egress	Yes	No	No	No
Layer 2 Bridge	Yes	Yes	Yes	No

All active RFC2544-based benchmarking tests are stopped when any of the following events takes place either in the initiator or in the reflector:

- System events such as Packet Forwarding Engine restarts, routing engine restarts, and so on.
- Test interface change events such as deactivation and reactivation of the interface, disabling and enabling of the interface, and so on.

After the benchmarking tests are stopped, the test states of the tests are removed and the user can restart the same test. Other ongoing tests on other interfaces are not interrupted.



**NOTE:** RFC2544-based benchmarking tests are not supported during unified in-service software upgrade (ISSU) and graceful routing engine switchover (GRES).

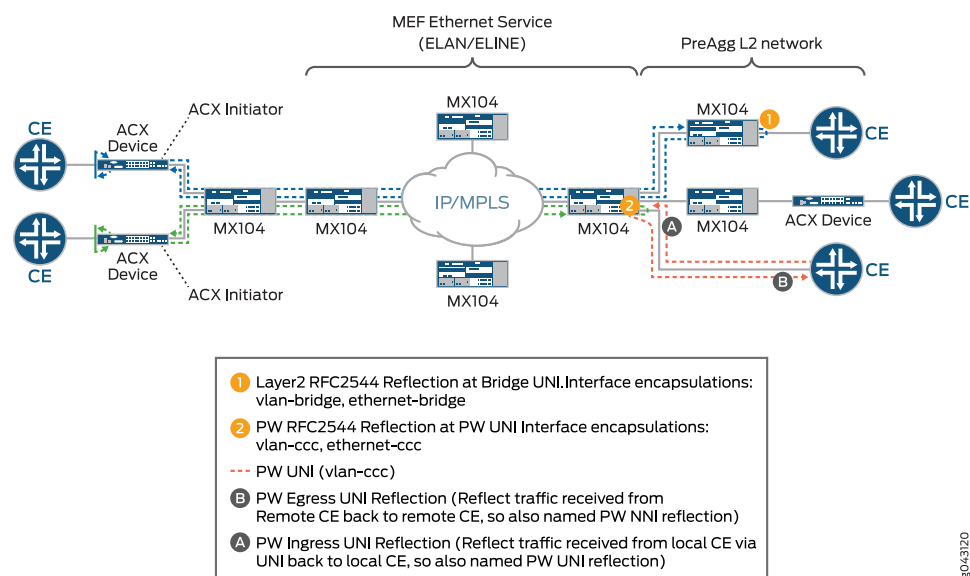
#### Related Documentation

- [Configuring an RFC 2544-Based Benchmarking Test on page 989](#)
- [Supported RFC2544-Based Benchmarking Statements on MX104 Routers on page 988](#)

## Layer 2 RFC2544-Based Benchmarking Tests Overview

The Metro Ethernet Forum (MEF) defines two Ethernet service types—E-LAN and E-Line—and specifies the associated service attributes and parameters. These services can be supported within the Metro Ethernet Network (MEN) and also supported over different transport technologies such as SONET, MPLS, and so on. Juniper networks ACX Series routers and MX104 routers provide support for Layer 2 E-LAN and E-Line services, pseudowire reflection, as well as IPv4 services. [Figure 36 on page 986](#) shows a sample topology for the E-LAN and E-Line reflection supported on MX104 routers.

**Figure 36: E-LAN and E-Line Reflection in Metro Solution**



MX104 routers support RFC2544-based benchmarking tests for Layer 2 reflection (E-LAN service) in basic bridge domains only. RFC2544-based benchmarking and performance measurement testing for Layer 2 services is supported on unicast traffic in egress direction only.

In an E-LAN service, during the benchmarking tests, the initiator or generator transmits a test packet (unicast) to a reflector. The reflector receives and reflects the test packet back to the initiator. The test packet is an UDP-over-IP packet with a source and destination MAC address. A Layer 2 traffic reflection session is identified by the source MAC address, the destination MAC address, and the egress interface. By default, RFC2544-based benchmarking tests are performed when there is no other service traffic. This mode of operation is known as out-of-service mode. The default service mode for the reflecting egress interface for an E-LAN service is also out-of-service mode. In out-of-service mode, while the test is running, all the data traffic sent to and from the UNI port under test on the service is interrupted. Control protocol peering is not interrupted whereas pass through control protocol packets such as end-to-end CFM sessions are interrupted. If you do not want the control protocol packets interrupted, you can configure the E-LAN service mode as in-service mode. In the in-service mode, while the test is

running, the rest of the data traffic flow sent to and from the UNI port under test on the service is not interrupted. Both peering and pass through control protocols are not interrupted.

By default, for E-LAN services, the default behavior of the reflector is to swap MAC addresses. The reflector swaps the source and destination MAC addresses and sends the packet back to the initiator. [Table 39 on page 987](#) describes the MAC address swapping behavior for the service types.

**Table 39: MAC Address Swapping Behavior for E-LAN and E-Line Services**

Family	Direction	Default Behavior	User-configurable
bridge	Egress	MAC address swap (E-LAN service type)	No
		No MAC address swap (E-Line service type)	Yes
ccc	Egress	No MAC address swap	No
	Ingress	MAC address swap	No

By default, the IP addresses and UDP ports are not modified. Optionally, you can configure the reflector to swap the source and destination IP address and the source and destination UDP ports.

You can configure an ACX Series router to operate as an initiator as well as a reflector. The MX104 router can be configured to operate only as a reflector.



**NOTE:** The maximum reflection bandwidth supported is 4Gbps. Because RFC2544 reflection shares system bandwidth with other loopback services such as tunnel services, you must manage the sharing of bandwidth for performing RFC2544-based performance tests.



**NOTE:** RFC2544-based benchmarking tests are not supported during unified in-service software upgrade (ISSU) and graceful routing engine switchover (GRES).

**Related Documentation**

- [RFC2544-Based Benchmarking Tests Overview on page 983](#)
- [Supported RFC2544-Based Benchmarking Statements on MX104 Routers on page 988](#)
- [Example: Configuring RFC2544-Based Benchmarking Tests for Layer 2 E-LAN Services in Bridge Domains on page 1016](#)

## Supported RFC2544-Based Benchmarking Statements on MX104 Routers

Table 40 on page 988 lists the reflector-specific configuration statements that are supported on the MX104 routers. Note that an (–) denotes that the command is not supported.

**Table 40: Supported RFC2544-Based Benchmarking Reflector Statements on MX104**

Statement	Options	Initial Release on MX104 Routers
<code>destination-ipv4-address</code>	–	13.3R1
<code>destination-mac-address</code>	–	14.2R1
<code>destination-udp-port</code>	–	13.3R1
<code>direction</code>	(egress   ingress)	13.3R1
<code>disable-signature-check</code>	–	15.1R1
<code>family</code>	(ccc   inet)	13.3R1
	(bridge   ccc   inet)	14.2R1
	(mpls   vpls)	15.1R1
<code>in-service</code>	–	14.2R1
<code>ip-swap</code>	–	14.2R1
<code>mode</code>	reflect	13.3R1
<code>reflect-etype</code>	–	15.1R1
<code>reflect-mode</code>	(mac-rewrite   mac-swap   no-mac-swap)	14.2R1
<code>service-type</code>	(eline   elan)	14.2R1
<code>source-ipv4-address</code>	–	13.3R1
<code>source-mac-address</code>	–	14.2R1
<code>source-udp-port</code>	–	13.3R1
<code>test-interface</code>	–	13.3R1
<code>udp-tcp-port-swap</code>	–	14.2R1

**Related** • [RFC2544-Based Benchmarking Tests Overview on page 983](#)

- Documentation**
- [Configuring an RFC 2544-Based Benchmarking Test on page 989](#)
  - [Example: Configuring RFC2544-Based Benchmarking Tests for Layer 2 E-LAN Services in Bridge Domains on page 1016](#)

## Configuring an RFC 2544-Based Benchmarking Test

You can configure a benchmarking test to detect and measure performance attributes, such as throughput, latency, frame loss, and bursty or back-to-back frames, of network devices. RFC 2544-based benchmarking test is performed by transmitting test packets from a device that functions as the generator or the initiator. These packets are sent to a device that functions as the reflector, which receives and returns the packets back to the initiator.



**NOTE:** The test configuration is applied only when you start the test. If you update the test configuration during the test, you have to start the test again for the updated configuration to take effect.

You must configure a test profile and reference the test profile in a unique test name that defines the parameters for the test to be performed on a certain device. However, the test profile is required when the test mode is configured as initiation and termination. The **test-profile** parameter is disregarded when the test mode is configured as reflection. MX104 routers support only the reflection function in the RFC 2544-based benchmarking tests. A reflection service does not use the parameters specified in the test profile.

The following topics describe how to configure a test name for an RFC 2544-based benchmarking test on an MX104 router for Layer 3 IPv4, Ethernet pseudowire, and Layer 2 bridge networks:

- [Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a IPv4 Network on page 989](#)
- [Configuring a Test Name for an RFC 2544-Based Benchmarking Test for an Ethernet Pseudowire on page 991](#)
- [Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a Layer 2 E-LAN Service in Bridge Domain on page 992](#)

### Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a IPv4 Network

You can configure a test name by including the **test-name test-name** statement at the **[edit services rpm rfc2544-benchmarking]** hierarchy level. In the test name, you can configure attributes of the test iteration, such as the address family (type of service, IPv4 or Ethernet), the logical interface, and test duration that are used for a benchmarking test to be run.

To configure a test name and define its attributes for an IPv4 network:

1. In configuration mode, go to the **[edit services]** hierarchy level.

**[edit]**

```
user@host# edit services
```

2. Configure a instance.

```
[edit services]
user@host# edit rpm
```

3. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

4. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

5. Specify the test mode for the packets that are sent during the benchmarking test. The **reflect** option causes the test frames to be reflected on the IPv4 network.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```

6. Configure the address type family for the benchmarking test. The **inet** option indicates that the test is run on an IPv4 service.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family inet
```

7. Configure the destination IPv4 address for the test packets. This parameter is required only if you configure IPv4 family **inet**. If you do not configure the destination IPv4 address, the default value of 192.168.1.20 is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set destination-ipv4-address address
```

8. Specify the UDP port of the destination to be used in the UDP header for the generated frames. If you do not specify the UDP port, the default value of 4041 is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set destination-udp-port port-number
```

9. (Optional) Specify the source IPv4 address to be used in generated test frames. If you do not configure the source IPv4 address for **inet** family, the source address of the interface is used to transmit the test frames.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set source-ipv4-address address
```

10. Specify the UDP port of the source to be used in the UDP header for the generated frames. If you do not specify the UDP port, the default value of 4041 is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set source-udp-port port-number
```

11. Specify the logical interface on which the RFC 2544-based benchmarking test is run. If you configure an **inet** family and the test mode to reflect the frames back on the sender from the other end, then the logical interface is used as the interface to enable the reflection service (reflection is performed on the packets entering the specified interface). If you not configure the logical interface for reflection test mode, then a



lookup is performed on the source IPv4 address to determine the interface that hosts the address.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface interface-name
```

## Configuring a Test Name for an RFC 2544-Based Benchmarking Test for an Ethernet Pseudowire

You can configure a test name by including the **test-name test-name** statement at the **[edit services rpm rfc2544-benchmarking]** hierarchy level. In the test name, you can configure attributes of the test iteration, such as the address family (type of service IPv4 or Ethernet), the logical interface, and test duration, that are used for a benchmarking test to be run. The test name combined with the test profile represent a single real-time performance monitoring (RPM) configuration instance.

To configure a test name and define its attributes for an Ethernet Pseudowire:

1. In configuration mode, go to the **[edit services]** hierarchy level.

```
[edit]
user@host# edit services
```

2. Configure an RPM service instance.

```
[edit services]
user@host# edit rpm
```

3. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

4. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

5. Specify the test mode for the packets that are sent during the benchmarking test. The **reflect** option causes the test frames to be reflected on the Ethernet pseudowire.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```

6. Configure the address type family for the benchmarking test. The **ccc** option indicates that the test is run on a CCC or Ethernet pseudowire service.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```

7. Specify the direction of the interface on which the test must be run. This parameter is valid only for a family. To enable the test to be run in the egress direction of the interface (network-to-network interface (NNI)), use the **egress** option. To enable the test to be run in the ingress direction of the interface (user-to-network interface (UNI)), use the **ingress** option.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction egress
```

8. (Optional) Specify the source IPv4 address to be used in generated test frames. If you do not configure the source IPv4 address for family, the default value of 192.168.1.10 is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set source-ipv4-address address
```

9. Specify the logical interface on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface interface-name
```

## Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a Layer 2 E-LAN Service in Bridge Domain

You can configure a test name by including the **test-name test-name** statement at the **[edit services rpm rfc2544-benchmarking]** hierarchy level. In the test name, you can configure attributes of the test iteration, such as the address family (bridge), the logical interface, and test duration, that are used for a benchmarking test to be run. The test name combined with the test profile represent a single real-time performance monitoring (RPM) configuration instance.

To configure a test name and define its attributes for a layer 2 E-LAN service in Bridge domains:

1. In configuration mode, go to the **[edit services]** hierarchy level.

```
[edit]
user@host# edit services
```

2. Configure an RPM service instance.

```
[edit services]
user@host# edit rpm
```

3. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

4. Define a name for the test—for example, l2b-test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name l2b-test1
```

5. Specify the source and destination MAC addresses of the test packet. Both these parameters are valid only for the bridge family.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set source-mac-address address destination-mac-address address
```

6. Specify the service type under test. This parameter is applicable only for the bridge family.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set service-type elan
```

7. Specify the test mode for the packets that are sent during the benchmarking test. The **reflect** option causes the test frames to be reflected over the Layer 2 bridge.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set mode reflect
```

8. Configure the address type family for the benchmarking test. The **bridge** option indicates that the test is run on a E-LAN service over a bridge domain.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set family bridge
```

9. Specify the direction of the interface on which the test must be run. This parameter is valid only for a family. To enable the test to be run in the egress direction of the interface (network-to-network interface (NNI)), use the **egress** option. To enable the test to be run in the ingress direction of the interface (user-to-network interface (UNI)), use the **ingress** option.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set direction egress
```

10. Specify the logical interface on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set test-interface interface-name
```

#### Related Documentation

- [RFC2544-Based Benchmarking Tests Overview on page 983](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for UNI Direction of Ethernet Pseudowires on page 1001](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for NNI Direction of Ethernet Pseudowires on page 1008](#)
- [Example: Configuring RFC2544-Based Benchmarking Tests for Layer 2 E-LAN Services in Bridge Domains on page 1016](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for Layer 3 IPv4 Services on page 993](#)

## Example: Configuring an RFC 2544-Based Benchmarking Test for Layer 3 IPv4 Services

- [Requirements on page 993](#)
- [Overview on page 994](#)
- [Configuration on page 994](#)
- [Verifying the Results of the Benchmarking Test for Layer 3 IPv4 Services on page 1000](#)

### Requirements

This example uses the following hardware and software components:

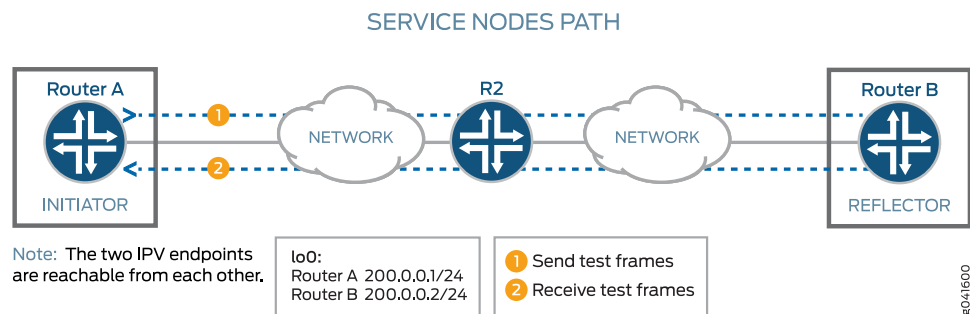
- 
- An ACX Series Universal Access Router—
- Junos OS Release or later

## Overview

Consider a sample topology in which a router, Router A, functions as an initiator and terminator of the test frames for an RFC 2544-based benchmarking test. Router A is connected over a Layer 3 network to another router, Router B, which functions as a reflector to reflect back the test frames it receives from Router A. IPv4 is used for transmission of test frames over the Layer 3 network. This benchmarking test is used to compute the IPv4 service parameters between Router A and Router B. Logical interfaces on both the routers are configured with IPv4 addresses to measure the performance attributes, such as throughput, latency, frame loss, and bursty frames, of network devices for the IPv4 service.

Figure 37 on page 994 shows the sample topology to perform an RFC 2544 test for a Layer 3 IPv4 service.

Figure 37: RFC 2544-Based Benchmarking Test for a Layer 3 IPv4 Service



## Configuration

In this example, you configure the benchmarking test for a Layer 3 IPv4 service that is between interface ge-0/0/0 on Router A and interface ge-0/0/4 on Router B to detect and analyze the performance of the interconnecting routers.

- [Configuring Benchmarking Test Parameters on Router A on page 995](#)
- [Configuring Benchmarking Test Parameters on Router B on page 997](#)
- [Results on page 999](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

### Configuring Benchmarking Test Parameters on Router A

```
set interfaces ge-0/0/0 unit 0 family inet address 200.0.0.1/24
set interfaces ge-0/0/0 unit 0 family mpls
set rfc2544-benchmarking profiles test-profile throughput test-type throughput
```

### Configuring Benchmarking Test Parameters on Router B

```
set rfc2544-benchmarking profiles test-profile throughput packet-size 64
set rfc2544-benchmarking profiles test-profile throughput test-duration 20m
set rfc2544-benchmarking profiles test-profile throughput bandwidth-kbps 500
set rfc2544-benchmarking tests test-name test1 test-profile throughput
set rfc2544-benchmarking tests test-name test1 interface ge-0/0/0.1
set rfc2544-benchmarking tests test-name test1 mode initiate-and-terminate
set rfc2544-benchmarking tests test-name test1 family inet
set rfc2544-benchmarking tests test-name test1 dest-address 200.0.0.2
set rfc2544-benchmarking tests test-name test1 udp-port 4001
```

```
set interfaces ge-0/0/4 unit 0 family inet address 200.0.0.2/24
set interfaces ge-0/0/4 unit 0 family mpls
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/4.1
set services rpm rfc2544-benchmarking tests test-name test1 mode reflect
set services rpm rfc2544-benchmarking tests test-name test1 family inet
set services rpm rfc2544-benchmarking tests test-name test1 dest-address 200.0.0.1
set services rpm rfc2544-benchmarking tests test-name test1 udp-port 4001
```

### Configuring Benchmarking Test Parameters on Router A

#### Step-by-Step Procedure

The following require you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the test parameters on Router A:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:  

```
[edit]
user@host# edit interfaces
```
2. Configure the interface on which the test must be run.  

```
[edit interfaces]
user@host# edit ge-0/0/0
```
3. Configure a logical unit and specify the protocol family.  

```
[edit interfaces ge-0/0/0]
user@host# edit unit 0 family inet
```
4. Specify the address for the logical interface.  

```
[edit interfaces ge-0/0/0 unit 0 family inet]
user@host# set address 200.0.0.1/24
```
5. Enter the **up** command to go the previous level in the configuration hierarchy.  

```
[edit interfaces ge-0/0/0 unit 0 family inet]
user@host# up
```
6. Configure the MPLS family on the logical interface.  

```
[edit interfaces ge-0/0/0 unit 0]
user@host# set family mpls
```

7. Go to the top level of the configuration command mode.  

```
[edit interfaces ge-0/0/0 unit 0]
user@host# top
```
8. In configuration mode, go to the **[edit services]** hierarchy level.  

```
[edit]
user@host# edit services
```
9. Configure a real-time performance monitoring service (RPM) instance.  

```
[edit services]
user@host# edit rpm
```
10. Configure an RFC 2544-based benchmarking test for the RPM instance.  

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```
11. Define a name for a test profile—for example, throughput.  

```
[edit services rpm rfc2544-benchmarking]
user@host# edit profiles test-profile throughput
```
12. Configure the type of test to be performed as throughput.  

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type throughput
```
13. Specify the size of the test packet as 64 bytes.  

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type packet-size 64
```
14. Specify the period for which the test is to be performed in hours, minutes, or seconds by specifying a number followed by the letter h (for hours), m (for minutes), or s (for seconds), respectively.  

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type test-duration 20m
```
15. Define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.  

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type bandwidth-kbps 500
```
16. Enter the **up** command to go the previous level in the configuration hierarchy.  

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# up
```
17. Enter the **up** command to go the previous level in the configuration hierarchy.  

```
[edit services rpm rfc2544-benchmarking profiles]
user@host# up
```
18. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.  

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

19. Specify the name of the test profile—for example, throughput—to be associated with a particular test name.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-profile throughput
```

20. Specify the logical interface, ge-0/0/0.1, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface ge-0/0/0.1
```

21. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode initiate-and-terminate
```

22. Configure the address type family, **inet**, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family inet
```

23. Configure the destination IPv4 address for the test packets.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set dest-address 200.0.0.2
```

24. Specify the UDP port of the destination to be used in the UDP header for the generated frames as 4001.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set udp-port 4001
```

25. Start the benchmarking test on the initiator.

```
user@> test services rpm rfc2544-benchmarking test test1 start
```

After the test is successfully completed, it is automatically stopped at the initiator.

---

### Configuring Benchmarking Test Parameters on Router B

#### Step-by-Step Procedure

The following you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the test parameters on Router B:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

```
[edit]
user@host# edit interfaces
```

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@host# edit ge-0/0/4
```

3. Configure a logical unit and specify the protocol family as **inet**.

```
[edit interfaces ge-0/0/4]
```

- user@host# edit unit 0 family inet
4. Specify the address for the logical interface.  
[edit interfaces ge-0/0/4 unit 0 family inet]  
user@host# set address 200.0.0.2/24
  5. Enter the **up** command to go the previous level in the configuration hierarchy.  
[edit interfaces ge-0/0/4 unit 0 family inet]  
user@host# up
  6. Configure the MPLS family on the logical interface.  
[edit interfaces ge-0/0/4 unit 0]  
user@host# set family mpls
  7. Go to the top level of the configuration command mode.  
[edit interfaces ge-0/0/4 unit 0]  
user@host# top
  8. In configuration mode, go to the **[edit services]** hierarchy level.  
[edit]  
user@host# edit services
  9. Configure a real-time performance monitoring service (RPM) instance.  
[edit services]  
user@host# edit rpm
  10. Configure an RFC 2544-based benchmarking test for the RPM instance.  
[edit services rpm]  
user@host# edit rfc2544-benchmarking
  11. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.  
[edit services rpm rfc2544-benchmarking]  
user@host# edit tests test-name test1
  12. Specify the logical interface, ge-0/0/4.1, on which the RFC 2544-based benchmarking test is run.  
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set test-interface ge-0/0/4.1
  13. Specify **reflect** as the test mode for the packets that are sent during the benchmarking test.  
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set mode reflect
  14. Configure the address type family, **inet**, for the benchmarking test.  
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set family inet
  15. Configure the destination IPv4 address for the test packets as 200.0.0.1.  
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set dest-address 200.0.0.1



16. Specify the UDP port of the destination to be used in the UDP header for the generated frames as 4001.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set udp-port 4001
```

17. Start the benchmarking test on the reflector.

```
user@host> test services rpm rfc2544-benchmarking test test1 start
```

After the test is successfully completed at the initiator, you can stop the test at the reflector by entering the `test services rpm rfc2544-benchmarking test test1` command.

## Results

In configuration mode, confirm your configuration on Router A and Router B by entering the `show` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Benchmarking Test Parameters on Router A:

```
[edit interfaces]
ge-0/0/0 {
 unit 0 {
 family inet {
 address 200.0.0.1/24;
 }
 family mpls;
 }
}

[edit services rpm]
rfc2544-benchmarking {
 profiles {
 test-profile throughput {
 test-type throughput
 packet-size 64;
 test-duration 20m;
 bandwidth-kbps 500;
 }
 }

 tests {
 test-name test1 {
 test-profile throughput;
 interface ge-0/0/0.1;
 mode initiate,terminate;
 family inet;
 dest-address 200.0.0.2
 udp-port 4001;
 }
 }
}
```

Benchmarking Test Parameters on Router B:

```
[edit interfaces]
ge-0/0/4 {
 unit 0 {
```

```
 family inet {
 address 200.0.0.2/24;
 }
 family mpls;
 }

[edit services rpm]
rfc2544-benchmarking {
 # Note, When in reflector mode, test profile is not needed
 tests {
 test-name test1 {
 interface ge-0/0/4.1;
 mode reflect;
 family inet;
 dest-address 200.0.0.1;
 udp-port 4001;
 }
 }
}
```

After you have configured the device, enter the **commit** command in configuration mode.

## Verifying the Results of the Benchmarking Test for Layer 3 IPv4 Services

Examine the results of the benchmarking test that is performed on the configured service between Router A and Router B.

- [Verifying the Benchmarking Test Results on page 1000](#)

---

### Verifying the Benchmarking Test Results

<b>Purpose</b>	Verify that the necessary and desired statistical values are displayed for the benchmarking test that is run on the configured service between Router A and Router B.
<b>Action</b>	In operational mode, enter the <b>show services rpm rfc2544-benchmarking (aborted-tests   active-tests   completed-tests   summary)</b> command to display information about the results of each category or state of the RFC 2544-based benchmarking test, such as aborted tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RFC2544-Based Benchmarking Tests Overview on page 983</a></li><li>• <a href="#">Configuring an RFC 2544-Based Benchmarking Test on page 989</a></li></ul>

## Example: Configuring an RFC 2544-Based Benchmarking Test for UNI Direction of Ethernet Pseudowires

---

This example shows how to configure the benchmarking test for the user-to-network interface (UNI) direction of an Ethernet pseudowire service.

- [Requirements on page 1001](#)
- [Overview on page 1001](#)
- [Configuration on page 1002](#)
- [Verifying the Results of the Benchmarking Test for UNI Direction of an Ethernet Pseudowire Service on page 1008](#)

### Requirements

This example uses the following hardware and software components:

- An ACX Series router—f
- Junos OS Release or later

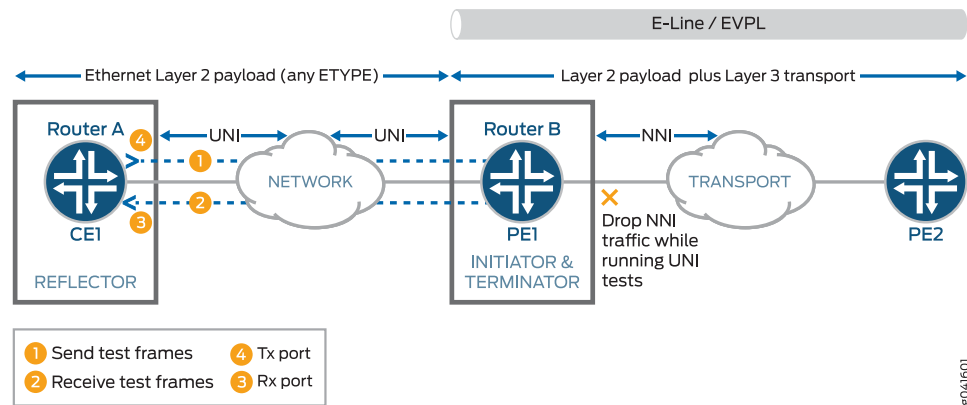
### Overview

Consider a sample topology in which a router, Router A, functions as a reflector of the test frames for an RFC 2544-based benchmarking test. The logical customer edge (CE)-facing interface and **inet** family are configured on Router A. Router A is not part of a pseudowire and therefore, a Layer 3 family configuration is required on it. Router A, which is a customer edge device CE1 is connected to Router B, which functions as a provider edge device PE1 over an Ethernet pseudowire in the UNI direction with EtherType or Layer 2 Ethernet payload. The logical interface, family, and UNI direction are configured on Router B. Router B or PE1 is connected over an Ethernet pseudowire in the NNI direction to a provider edge device at the remote site, PE2. The link between CE1 and PE1 is an Ethernet Layer 2 network and it can be configured with any EtherType value. The link between PE1 and PE2 is an Ethernet line (E-LINE) or an Ethernet Private Line (EPL) that has Layer 2 payload and Layer 3 transport sent over it. Router B or PE1 functions as an initiator and terminator of the test frames that are sent to Router A and reflected back from it.

This benchmarking test is used to compute the performance attributes in the user-to-network interface (UNI) direction of an Ethernet pseudowire service between Router A and Router B. Data traffic arriving from a network-to-network interface (NNI) toward the customer edge is ignored while the test is in progress. Packets from the CE are not sent toward the NNI because all packets are assumed to be test probes.

[Figure 38 on page 1002](#) shows the sample topology to perform an RFC 2544 test for the UNI direction of an Ethernet pseudowire service.

Figure 38: RFC 2544-Based Benchmarking Test for UNI Direction of an Ethernet Pseudowire



## Configuration

In this example, you configure the benchmarking test for the UNI direction of an Ethernet pseudowire service that is enabled between two routers to detect and analyze the performance of the interconnecting routers.

- [Configuring Benchmarking Test Parameters on Router A on page 1003](#)
- [Configuring Benchmarking Test Parameters on Router B on page 1005](#)
- [Results on page 1007](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

### Configuring Benchmarking Test Parameters on Router A

```
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 0 vlan-id 101
set interfaces ge-0/0/0 unit 0 family inet address 200.0.0.1/24
set services rpm rfc2544-benchmarking profiles test-profile throughput test-type throughput
set services rpm rfc2544-benchmarking profiles test-profile throughput packet-size 64
set services rpm rfc2544-benchmarking profiles test-profile throughput test-duration 20m
set services rpm rfc2544-benchmarking profiles test-profile throughput bandwidth-kbps 500
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/0.1
set services rpm rfc2544-benchmarking tests test-name test1 test-profile throughput
set services rpm rfc2544-benchmarking tests test-name test1 mode initiate,terminate
set services rpm rfc2544-benchmarking tests test-name test1 family inet
set services rpm rfc2544-benchmarking tests test-name test1 dest-address 200.0.0.2
set services rpm rfc2544-benchmarking tests test-name test1 udp-port 4001
```

## Configuring Benchmarking Test Parameters on Router B

```
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 0 encapsulation vlan-ccc
set interfaces ge-0/0/4 unit 0 vlan-id 101
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/4.1
set services rpm rfc2544-benchmarking tests test-name test1 mode reflect
set services rpm rfc2544-benchmarking tests test-name test1 mode family ccc
set services rpm rfc2544-benchmarking tests test-name test1 direction uni
```

---

### Configuring Benchmarking Test Parameters on Router A

#### Step-by-Step Procedure

The following require you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the test parameters on Router A:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:  

```
[edit]
user@host# edit interfaces
```
2. Configure the interface on which the test must be run.  

```
[edit interfaces]
user@host# edit ge-0/0/0
```
3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.  

```
[edit interfaces ge-0/0/0]
user@host# set vlan-tagging
```
4. Configure a logical unit and specify the protocol family as **inet**.  

```
[edit interfaces ge-0/0/0]
user@host# edit unit 0 family inet
```
5. Specify the address for the logical interface.  

```
[edit interfaces ge-0/0/0 unit 0 family inet]
user@host# set address 200.0.0.1/24
```
6. Configure the VLAN ID on the logical interface as 101.  

```
[edit interfaces ge-0/0/0 unit 0]
user@host# set vlan-id 101
```
7. Go to the top level of the configuration command mode.  

```
[edit interfaces ge-0/0/0 unit 0]
user@host# top
```
8. In configuration mode, go to the **[edit services]** hierarchy level.  

```
[edit]
user@host# edit services
```

9. Configure a real-time performance monitoring service (RPM) instance.  
[edit services]  
user@host# **edit rpm**
10. Configure an RFC 2544-based benchmarking test for the RPM instance.  
[edit services rpm]  
user@host# **edit rfc2544-benchmarking**
11. Define a name for a test profile—for example, throughput.  
[edit services rpm rfc2544-benchmarking]  
user@host# **edit profiles test-profile throughput**
12. Configure the type of test to be performed as throughput.  
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# **set test-type throughput**
13. Specify the size of the test packet as 64 bytes.  
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# **set test-type packet-size 64**
14. Specify the period for which the test is to be performed in hours, minutes, or seconds by specifying a number followed by the letter h (for hours), m (for minutes), or s (for seconds). In this example, you configure the period as 20 minutes.  
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# **set test-type test-duration 20m**
15. Define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.  
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# **set test-type bandwidth-kbps 500**
16. Enter the **up** command to go the previous level in the configuration hierarchy.  
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# **up**
17. Enter the **up** command to go the previous level in the configuration hierarchy.  
[edit services rpm rfc2544-benchmarking profiles]  
user@host# **up**
18. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.  
[edit services rpm rfc2544-benchmarking]  
user@host# **edit tests test-name test1**
19. Specify the name of the test profile—for example, throughput—to be associated with a particular test name.  
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# **set test-profile throughput**
20. Specify the logical interface, ge-0/0/0.1, on which the RFC 2544-based benchmarking test is run.  
[edit services rpm rfc2544-benchmarking tests test-name test1]

```
user@host# set test-interface ge-0/0/0.1
```

21. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode initiate-and-terminate
```

22. Configure the address type family, **inet**, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family inet
```

23. Configure the destination IPv4 address for the test packets as 200.0.0.2.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set dest-address 200.0.0.2
```

24. Specify the UDP port of the destination to be used in the UDP header for the generated frames as 4001.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set udp-port 4001
```

---

### Configuring Benchmarking Test Parameters on Router B

#### Step-by-Step Procedure

The following require you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the test parameters on Router B:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

```
[edit]
user@host# edit interfaces
```

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@host# edit ge-0/0/4
```

3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.

```
[edit interfaces ge-0/0/4]
user@host# set vlan-tagging
```

4. Configure a logical unit for the interface.

```
[edit interfaces ge-0/0/4]
user@host# edit unit 0
```

5. Specify the encapsulation for Ethernet VLAN circuits.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# set encapsulation vlan-ccc
```

6. Configure the VLAN ID as 101 on the logical interface.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# set vlan-id 101
```

7. Go to the top level of the configuration command mode.  

```
[edit interfaces ge-0/0/4 unit 0]
user@host# top
```
8. In configuration mode, go to the **[edit services]** hierarchy level.  

```
[edit]
user@host# edit services
```
9. Configure a real-time performance monitoring service (RPM) instance.  

```
[edit services]
user@host# edit rpm
```
10. Configure an RFC 2544-based benchmarking test for the RPM instance.  

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```
11. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.  

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```
12. Specify the logical interface on which the RFC 2544-based benchmarking test is run.  

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface ge-0/0/4.1
```
13. Specify **reflect** as the test mode for the packets that are sent during the benchmarking test.  

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```
14. Configure the address type family, **ccc**, for the benchmarking test.  

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```
15. Specify the direction of the interface on which the test must be run, which is UNI in this example.  

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction uni
```
16. Start the benchmarking test on the reflector.  

```
user@host> test services rpm rfc2544-benchmarking test test1 start
```

After the test is successfully completed at the initiator, you can stop the test at the reflector by entering the **test services rpm rfc2544-benchmarking test test1 stop** command.



## Results

In configuration mode, confirm your configuration on Router A and Router B by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Benchmarking Test Parameters on Router A:

```
[edit interfaces]
ge-0/0/0 {
 vlan-tagging;
 unit 0 {
 vlan-id 101;
 family inet {
 address 200.0.0.1/24;
 }
 }
}

[edit services rpm]
rfc2544-benchmarking {
 profiles {
 test-profile throughput {
 test-type throughput
 packet-size 64;
 test-duration 20m;
 bandwidth-kbps 500;
 }
 }

 tests {
 test-name test1 {
 interface ge-0/0/0.1;
 test-profile throughput;
 mode initiate,terminate;
 family inet;
 dest-address 200.0.0.2
 udp-port 4001;
 }
 }
}
```

Benchmarking Test Parameters on Router B:

```
[edit interfaces]
ge-0/0/4 {
 vlan-tagging;
 unit 0 {
 encapsulation vlan-ccc;
 vlan-id 101;
 }
}

[edit services rpm]
rfc2544-benchmarking {
 # Note, When in reflector mode, test profile is not needed
 tests {
 test-name test1 {
 interface ge-0/0/4.1;
```

```
 mode reflect;
 family ccc;
 direction uni;
 }
}
```

After you have configured the device, enter the **commit** command in configuration mode.

## Verifying the Results of the Benchmarking Test for UNI Direction of an Ethernet Pseudowire Service

Examine the results of the benchmarking test that is performed on the configured service between Router A and Router B.

- [Verifying the Benchmarking Test Results on page 1008](#)

---

### Verifying the Benchmarking Test Results

<b>Purpose</b>	Verify that the necessary and desired statistical values are displayed for the benchmarking test that is run on the configured service between Router A and Router B.
<b>Action</b>	In operational mode, enter the <b>show services rpm rfc2544-benchmarking (aborted-tests   active-tests   completed-tests   summary)</b> command to display information about the results of each category or state of the RFC 2544-based benchmarking test, such as aborted tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance.
<b>Meaning</b>	The output displays the details of the benchmarking test that was performed. For more information about the <b>show services rpm rfc2544-benchmarking</b> operational command, see <b>show services rpm rfc2544-benchmarking</b> in the <a href="#">CLI Explorer</a> .
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RFC2544-Based Benchmarking Tests Overview on page 983</a></li><li>• <a href="#">Configuring an RFC 2544-Based Benchmarking Test on page 989</a></li></ul>

---

## Example: Configuring an RFC 2544-Based Benchmarking Test for NNI Direction of Ethernet Pseudowires

This example shows how to configure the benchmarking test for a network-to-network interface (NNI) direction of an Ethernet pseudowire service.

- [Requirements on page 1009](#)
- [Overview on page 1009](#)
- [Configuration on page 1010](#)
- [Verifying the Results of the Benchmarking Test for NNI Direction of an Ethernet Pseudowire Service on page 1016](#)

## Requirements

This example uses the following hardware and software components:

- An ACX Series router
- Junos OS Release or later

## Overview

Consider a sample topology in which a router, Router A, functions as an initiator and terminator of the test frames for an RFC 2544-based benchmarking test. Router A operates as a provider edge device PE1, which is connected to a customer edge device CE1 on one side and over an Ethernet pseudowire to another router Router B, which functions as a reflector to reflect back the test frames it receives from Router A. Router B operates as a provider edge device, PE2, which is the remote router located at the other side of the service provider core. The UNI direction of CE1 is connected to the NNI direction of PE1. An MPLS tunnel connects PE1 and PE2 over the Ethernet pseudowire or the Ethernet line (E-LINE).

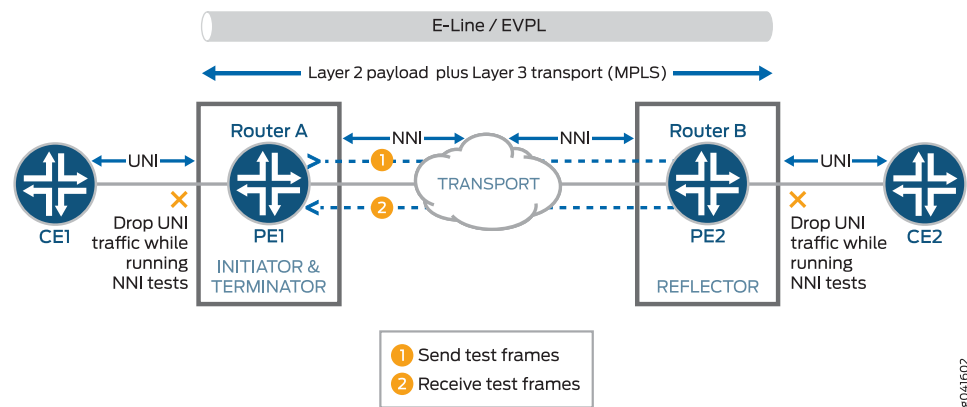


**NOTE:** When pseudowire reflection is enabled on an interface, the router does not block the ingress or egress traffic through the test interface. To block other data traffic, you must explicitly configure firewall filters.

This benchmarking test is used to compute the performance attributes in the network-to-network interface (NNI) direction of an Ethernet pseudowire service between Router A and Router B. The logical interface under test on Router A is the CE1 interface with UNI as the direction, and the logical interface under test on Router B is the CE2 interface with NNI as the direction. Data traffic arriving from UNI toward NNI is ignored while the test is in progress. Packets from NNI are not sent toward the customer edge because all packets are assumed to be test frames. The family and NNI direction are configured on routers A and B.

[Figure 39 on page 1010](#) shows the sample topology to perform an RFC 2544 test for the NNI direction of an Ethernet pseudowire service.

Figure 39: RFC 2544-Based Benchmarking Test for NNI Direction of an Ethernet Pseudowire



## Configuration

In this example, you configure the benchmarking test for the NNI direction of an Ethernet pseudowire service that is enabled between two routers to detect and analyze the performance of the interconnecting routers.

- [Configuring Benchmarking Test Parameters on Router A on page 1011](#)
- [Configuring Benchmarking Test Parameters on Router B on page 1013](#)
- [Results on page 1015](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

### Configuring Benchmarking Test Parameters on Router A

```
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 0 encapsulation vlan-ccc
set interfaces ge-0/0/0 unit 0 vlan-id 101
set services rpm rfc2544-benchmarking profiles test-profile throughput test-type throughput
set services rpm rfc2544-benchmarking profiles test-profile throughput packet-size 64
set services rpm rfc2544-benchmarking profiles test-profile throughput test-duration 20
set services rpm rfc2544-benchmarking profiles test-profile throughput bandwidth-kbps 500
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/0.1
set services rpm rfc2544-benchmarking tests test-name test1 test-profile throughput
set services rpm rfc2544-benchmarking tests test-name test1 mode initiate,terminate
set services rpm rfc2544-benchmarking tests test-name test1 family ccc
set services rpm rfc2544-benchmarking tests test-name test1 direction nni
```

### Configuring Benchmarking Test

**Parameters on Router****B**

```

set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 0 encapsulation vlan-ccc
set interfaces ge-0/0/4 unit 0 vlan-id 101
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/4.1
set services rpm rfc2544-benchmarking tests test-name test1 mode reflect
set services rpm rfc2544-benchmarking tests test-name test1 mode family ccc
set services rpm rfc2544-benchmarking tests test-name test1 direction uni

```

**Configuring Benchmarking Test Parameters on Router****Step-by-Step  
Procedure**

The following require you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the test parameters on Router A:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:
 

```

[edit]
user@host# edit interfaces

```
2. Configure the interface on which the test must be run.
 

```

[edit interfaces]
user@host# edit ge-0/0/0

```
3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.
 

```

[edit interfaces ge-0/0/0]
user@host# set vlan-tagging

```
4. Configure a logical unit for the interface.
 

```

[edit interfaces ge-0/0/0]
user@host# edit unit 0

```
5. Specify the encapsulation for Ethernet VLAN circuits.
 

```

[edit interfaces ge-0/0/0 unit 0]
user@host# set encapsulation vlan-ccc

```
6. Configure the VLAN ID on the logical interface.
 

```

[edit interfaces ge-0/0/0 unit 0]
user@host# set vlan-id 101

```
7. Go to the top level of the configuration command mode.
 

```

[edit interfaces ge-0/0/0 unit 0]
user@host# top

```
8. In configuration mode, go to the **[edit services]** hierarchy level.
 

```

[edit]
user@host# edit services

```
9. Configure a real-time performance monitoring service (RPM) instance.

- ```
[edit services]
user@host# edit rpm
```
10. Configure an RFC 2544-based benchmarking test for the RPM instance.
- ```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```
11. Define a name for a test profile—for example, throughput.
- ```
[edit services rpm rfc2544-benchmarking]
user@host# edit profiles test-profile throughput
```
12. Configure the type of test to be performed as throughput.
- ```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type throughput
```
13. Specify the size of the test packet as 64 bytes.
- ```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type packet-size 64
```
14. Specify the period—for example, 20 minutes—for which the test is to be performed in hours, minutes, or seconds by specifying a number followed by the letter h (for hours), m (for minutes), or s (for seconds).
- ```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type test-duration 20m
```
15. Define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.
- ```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type bandwidth-kbps 500
```
16. Enter the **up** command to go the previous level in the configuration hierarchy.
- ```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# up
```
17. Enter the **up** command to go the previous level in the configuration hierarchy.
- ```
[edit services rpm rfc2544-benchmarking profiles]
user@host# up
```
18. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.
- ```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```
19. Specify the name of the test profile—for example, throughput—to be associated with a particular test name.
- ```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-profile throughput
```
20. Specify the logical interface, ge-0/0/0.1, on which the RFC 2544-based benchmarking test is run.
- ```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface ge-0/0/0.1
```

21. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode initiate-and-terminate
```

22. Configure the address type family, **ccc**, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```

23. Specify the direction of the interface on which the test must be run, which is **NNI** in this example.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction nni
```

---

### Configuring Benchmarking Test Parameters on Router B

#### Step-by-Step Procedure

The following require you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the test parameters on Router B:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

```
[edit]
user@host# edit interfaces
```

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@host# edit ge-0/0/4
```

3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.

```
[edit interfaces ge-0/0/4]
user@host# set vlan-tagging
```

4. Configure a logical unit for the interface.

```
[edit interfaces ge-0/0/4]
user@host# edit unit 0
```

5. Specify the encapsulation for Ethernet VLAN circuits.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# set encapsulation vlan-ccc
```

6. Configure the VLAN ID on the logical interface.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# set vlan-id 101
```

7. Go to the top level of the configuration command mode.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# top
```

8. In configuration mode, go to the **[edit services]** hierarchy level.

```
[edit]
user@host# edit services
```

9. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]
user@host# edit rpm
```

10. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

11. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

12. Specify the logical interface, ge-0/0/4.1, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface ge-0/0/4.1
```



**NOTE:** When pseudowire reflection is enabled on an interface, the router does not block the ingress or egress traffic through the test interface. To block other data traffic, you must explicitly configure firewall filters.

13. Specify **reflect** as the test mode for the packets that are sent during the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```

14. Configure the address type family, **ccc**, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```

15. Specify the direction of the interface on which the test must be run, which is **NNI** in this example.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction nni
```

16. Start the benchmarking test on the reflector.

```
user@host> test services rpm rfc2544-benchmarking test test1 start
```

After the test is successfully completed at the initiator, you can stop the test at the reflector by entering the **test services rpm rfc2544-benchmarking test test1 stop** command.



## Results

In configuration mode, confirm your configuration on Router A and Router B by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Benchmarking Test Parameters on Router A:

```
[edit interfaces]
ge-0/0/0 {
 vlan-tagging;
 unit 0 {
 encapsulation vlan-ccc;
 vlan-id 101;
 }
}

[edit services rpm]
rfc2544-benchmarking {
 profiles {
 test-profile throughput {
 test-type throughput
 packet-size 64;
 test-duration 20m;
 bandwidth-kbps 500;
 }
 }

 tests {
 test-name test1 {
 interface ge-0/0/0.1;
 test-profile throughput;
 mode initiate,terminate;
 family ccc;
 direction nni;
 }
 }
}
```

Benchmarking Test Parameters on Router B:

```
[edit interfaces]
ge-0/0/4 {
 vlan-tagging;
 unit 0 {
 encapsulation vlan-ccc;
 vlan-id 101;
 }
}

[edit services rpm]
rfc2544-benchmarking {
 # Note, When in reflector mode, test profile is not needed
 tests {
 test-name test1 {
 interface ge-0/0/4.1;
 mode reflect;
 family ccc;
 }
 }
}
```

```
 direction nni;
 }
}
```

After you have configured the device, enter the **commit** command in configuration mode.

## Verifying the Results of the Benchmarking Test for NNI Direction of an Ethernet Pseudowire Service

Examine the results of the benchmarking test that is performed on the configured service between Router A and Router B.

- [Verifying the Benchmarking Test Results on page 1016](#)

---

### Verifying the Benchmarking Test Results

<b>Purpose</b>	Verify that the necessary and desired statistical values are displayed for the benchmarking test that is run on the configured service between Router A and Router B.
<b>Action</b>	In operational mode, enter the <b>show services rpm rfc2544-benchmarking (aborted-tests   active-tests   completed-tests   summary)</b> command to display information about the results of each category or state of the RFC 2544-based benchmarking test, such as aborted tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance.
<b>Meaning</b>	The output displays the details of the benchmarking test that was performed. For more information about the <b>show services rpm rfc2544-benchmarking</b> operational command, see <b>show services rpm rfc2544-benchmarking</b> in the <a href="#">CLI Explorer</a> .
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RFC2544-Based Benchmarking Tests Overview on page 983</a></li><li>• <a href="#">Configuring an RFC 2544-Based Benchmarking Test on page 989</a></li></ul>

---

## Example: Configuring RFC2544-Based Benchmarking Tests for Layer 2 E-LAN Services in Bridge Domains

This example shows how to configure benchmarking tests for the Layer 2 E-LAN services in bridge domains. The example covers the four basic tests: throughput, frame-loss, back-to-back, and latency.

- [Requirements on page 1017](#)
- [Overview on page 1017](#)
- [Configuration on page 1018](#)
- [Verifying the Results of the Benchmarking Tests for Layer 2 Services \(E-LAN\) in Bridge Domains on page 1032](#)

## Requirements

This example uses the following hardware and software components:

- An MX104 3D Universal Edge router
- An ACX Series router
- Junos OS Release 14.2 or later for MX Series routers

## Overview


Consider a sample topology in which an ACX router functions as an initiator and terminator of the test frames for an RFC2544-based benchmarking test. ACX router is connected to a customer edge device CE1, on one side and is connected over a Layer 2 network to an MX104 router. The MX104 router functions as a reflector to reflect the test frames it receives from the ACX Series initiator back to the initiator. The MX04 router is also connected to a customer edge device CE2.



**NOTE:** When Layer 2 reflection is enabled on an interface, filters are configured internally to block the ingress and egress traffic except test traffic through the test interface.

Figure 40 on page 1017 shows the sample topology to perform all four RFC2544-based benchmarking tests (throughput, back-to-back frames, latency, and frame-loss) for the UNI direction on a Layer 2 bridge network.

**Figure 40: Layer 2 reflection Simple Topology**



On the ACX router, ge-1/2/1.0 is the Layer 2 NNI interface and ge-1/1/3.0 is the Layer 2 UNI interface. On the MX104 router, ge-1/1/6.0 is the Layer 2 NNI interface and ge-1/1/5.0 is the Layer 2 UNI interface. The benchmarking tests are used to compute the performance attributes for an E-LAN service on a bridge domain.



**NOTE:** Test packets can be identified using the destination MAC address, source MAC address, and test interface. Both tagged and untagged interfaces are supported. For tagged interfaces, the test interface is the VLAN sub interface. For untagged interfaces, the physical port represents the test interface. Traffic through other VLAN sub interfaces, present in the same physical port, is not affected when you configure the benchmarking test on one of the sub interfaces.

## Configuration

In this example, you configure the benchmarking tests for the UNI direction for an E-LAN service on a Layer 2 bridge domain that is enabled between two routers to detect and analyze the performance of the interconnected routers. In this example, we start by configuring the ACX Series router. On the ACX router, you first configure each test by specifying the test profile, the test attributes, and then define the test by associating the test with the test profile with the relevant attributes. You can then configure the interface. On the MX104 router, you will perform the same steps. However, a few attributes such as the outer VLAN ID, source UDP port, destination UDP port, the duration of each iteration, and their values are only applicable to the initiator or the ACX router.



**NOTE:** When you configure the Layer 2 reflection, you can specify the service type under test as ELINE if you want to simulate an ELINE service using bridge encapsulation.

- [Configuring Throughput Benchmarking Test Parameters on the ACX Series Router on page 1021](#)
- [Configuring Back-to-Back Frames Benchmarking Test Parameters on the ACX Series Router on page 1022](#)
- [Configuring Latency Benchmarking Test Parameters on the ACX Series Router on page 1023](#)
- [Configuring Frame Loss Benchmarking Test Parameters on the ACX Series Router on page 1025](#)
- [Configuring Other Benchmarking Test Parameters on the ACX Series Router on page 1026](#)
- [Configuring Benchmarking Test Parameters on the MX104 Router on page 1027](#)
- [Configuring Other Benchmarking Test Parameters on the MX104 Router on page 1028](#)
- [Results on page 1029](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

### Configuring Benchmarking Test

```
set services rpm rfc2544-benchmarking profiles test-profile tput test-type throughput
set services rpm rfc2544-benchmarking profiles test-profile tput packet-size 128
set services rpm rfc2544-benchmarking profiles test-profile tput bandwidth-kbps 900000
```

## Parameters on the ACX Series Router

```

set services rpm rfc2544-benchmarking profiles test-profile b2bt test-type
back-back-frames
set services rpm rfc2544-benchmarking profiles test-profile b2bt packet-size 512
set services rpm rfc2544-benchmarking profiles test-profile b2bt bandwidth-kbps 950000
set services rpm rfc2544-benchmarking profiles test-profile lty test-type latency
set services rpm rfc2544-benchmarking profiles test-profile lty packet-size 512
set services rpm rfc2544-benchmarking profiles test-profile lty bandwidth-kbps 1000000
set services rpm rfc2544-benchmarking profiles test-profile frloss test-type frame-loss
set services rpm rfc2544-benchmarking profiles test-profile frloss packet-size 1600
set services rpm rfc2544-benchmarking profiles test-profile frloss bandwidth-kbps
1000000
set services rpm rfc2544-benchmarking tests test-name tput-test test-profile tput
set services rpm rfc2544-benchmarking tests test-name tput-test source-mac-address
00:00:00:00:11:11
set services rpm rfc2544-benchmarking tests test-name tput-test destination-mac-address
00:00:00:00:22:22
set services rpm rfc2544-benchmarking tests test-name tput-test ovlan-id 400
set services rpm rfc2544-benchmarking tests test-name tput-test service-type elan
set services rpm rfc2544-benchmarking tests test-name tput-test mode
initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name tput-test family bridge
set services rpm rfc2544-benchmarking tests test-name tput-test direction egress
set services rpm rfc2544-benchmarking tests test-name tput-test source-udp-port 200
set services rpm rfc2544-benchmarking tests test-name tput-test destination-udp-port
200
set services rpm rfc2544-benchmarking tests test-name tput-test test-iterator-duration
20
set services rpm rfc2544-benchmarking tests test-name tput-test test-interface ge-1/1/3.0
set services rpm rfc2544-benchmarking tests test-name b2b-test test-profile b2bt
set services rpm rfc2544-benchmarking tests test-name b2b-test source-mac-address
00:00:00:00:11:11
set services rpm rfc2544-benchmarking tests test-name b2b-test destination-mac-address
00:00:00:00:22:22
set services rpm rfc2544-benchmarking tests test-name b2b-test ovlan-id 400
set services rpm rfc2544-benchmarking tests test-name b2b-test service-type elan
set services rpm rfc2544-benchmarking tests test-name b2b-test mode
initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name b2b-test family bridge
set services rpm rfc2544-benchmarking tests test-name b2b-test direction egress
set services rpm rfc2544-benchmarking tests test-name b2b-test test-iterator-duration
20
set services rpm rfc2544-benchmarking tests test-name b2b-test test-interface ge-1/1/3.0
set services rpm rfc2544-benchmarking tests test-name lty-test test-profile lty
set services rpm rfc2544-benchmarking tests test-name lty-test source-mac-address
00:00:00:00:11:11
set services rpm rfc2544-benchmarking tests test-name lty-test destination-mac-address
00:00:00:00:22:22
set services rpm rfc2544-benchmarking tests test-name lty-test ovlan-id 400
set services rpm rfc2544-benchmarking tests test-name lty-test service-type elan
set services rpm rfc2544-benchmarking tests test-name lty-test mode
initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name lty-test family bridge
set services rpm rfc2544-benchmarking tests test-name lty-test direction egress
set services rpm rfc2544-benchmarking tests test-name lty-test source-udp-port 200
set services rpm rfc2544-benchmarking tests test-name lty-test destination-udp-port
200

```

```
set services rpm rfc2544-benchmarking tests test-name lty-test test-iterator-duration 20
set services rpm rfc2544-benchmarking tests test-name lty-test test-interface ge-1/1/3.0
set services rpm rfc2544-benchmarking tests test-name frloss-test test-profile frloss
set services rpm rfc2544-benchmarking tests test-name frloss-test source-mac-address
 00:00:00:00:11:11
set services rpm rfc2544-benchmarking tests test-name frloss-test
 destination-mac-address 00:00:00:00:22:22
set services rpm rfc2544-benchmarking tests test-name frloss-test ovlan-id 400
set services rpm rfc2544-benchmarking tests test-name frloss-test service-type elan
set services rpm rfc2544-benchmarking tests test-name frloss-test mode
 initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name frloss-test family bridge
set services rpm rfc2544-benchmarking tests test-name frloss-test direction egress
set services rpm rfc2544-benchmarking tests test-name frloss-test source-udp-port 200
set services rpm rfc2544-benchmarking tests test-name frloss-test destination-udp-port
 200
set services rpm rfc2544-benchmarking tests test-name frloss-test test-iterator-duration
 20
set services rpm rfc2544-benchmarking tests test-name frloss-test test-interface ge-1/1/3.0
set interfaces ge-1/2/1 flexible-vlan-tagging
set interfaces ge-1/2/1 mtu 9192
set interfaces ge-1/2/1 encapsulation flexible-ethernet-services
set interfaces ge-1/2/1 unit 0 encapsulation vlan-bridge
set interfaces ge-1/2/1 unit 0 vlan-id 400
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 mtu 9192
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/3 unit 0 vlan-id 400
set bridge-domains bd1 vlan-id 600
set bridge-domains bd1 interface ge-1/2/1.0
set bridge-domains bd1 interface ge-1/1/3.0
```

**Configuring  
Benchmarking Test  
Parameters on the  
MX104 Router**

```
set services rpm rfc2544-benchmarking tests test-name l2b-reflector source-mac-address
 00:00:00:00:11:11
set services rpm rfc2544-benchmarking tests test-name l2b-reflector
 destination-mac-address 00:00:00:00:22:22
set services rpm rfc2544-benchmarking tests test-name l2b-reflector service-type elan
set services rpm rfc2544-benchmarking tests test-name l2b-reflector mode reflect
set services rpm rfc2544-benchmarking tests test-name l2b-reflector family bridge
set services rpm rfc2544-benchmarking tests test-name l2b-reflector direction egress
set services rpm rfc2544-benchmarking tests test-name l2b-reflector test-interface
 ge-1/1/5.0
set interfaces ge-1/1/6 flexible-vlan-tagging
set interfaces ge-1/1/6 mtu 9192
set interfaces ge-1/1/6 encapsulation flexible-ethernet-services
set interfaces ge-1/1/6 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/6 unit 0 vlan-id 100
set interfaces ge-1/1/5 flexible-vlan-tagging
set interfaces ge-1/1/5 mtu 9192
set interfaces ge-1/1/5 encapsulation flexible-ethernet-services
set interfaces ge-1/1/5 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/5 unit 0 vlan-id 100
set bridge-domains bd1 domain-type bridge
set bridge-domains bd1 vlan-id 500
set bridge-domains bd1 interface ge-1/1/6.0
```

```
set bridge-domains bd1 interface ge-1/1/5.0
```

### Configuring Throughput Benchmarking Test Parameters on the ACX Series Router

#### Step-by-Step Procedure

The following configuration requires you to configure a test profile for the throughput test and reference the test-profile in a unique test-name. The test-name defines the parameters for the throughput test to be performed on the ACX router.

To configure the throughput test parameters on the ACX Router:

1. In configuration mode, at the `[edit]` hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```

2. Define a name for the first test profile—for example, `tput` for the throughput test profile.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit profiles test-profile tput
```

3. Configure the type of test to be performed as throughput, specify the packet size as 128 bytes, and define the theoretical maximum bandwidth for the test in kilobits per second (Kbps), with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles test-profile tput]
user@host# set test-type throughput packet-size 128 bandwidth-kbps 900000
```

4. Enter the `up` command twice to go to the `[edit services rpm rfc2544-benchmarking]` level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile tput]
user@host# up
user@host# up
```

5. Define a name for the throughput test—for example, `tput-test`. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name tput-test
```

6. Specify the name of the test profile, `tput`, to be associated with the test name.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set test-profile tput
```

7. Configure the source and destination MAC address for the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set source-mac-address 00:00:00:00:11:11 destination-mac-address
00:00:00:00:22:22
```

8. Configure the outer VLAN ID for the test frames and specify the service type under test to be E-LAN.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set ovlan-id 400 service-type elan
```

9. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set mode initiate-and-terminate
```

10. Configure the family type, **bridge**, for the benchmarking test and specify the direction, egress. Also, specify the source and destination UDP port to be used in the UDP headers of the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set family bridge direction egress source-udp-port 200
destination-udp-port 200
```

11. Specify the duration of each iteration in seconds, with a value from 10 seconds to 1,728,000 seconds, and specify the logical interface, ge-0/2/1.0, on which the RFC2544-benchmarking tests are run.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set test-iterator-duration 20 test-interface ge-1/1/3.0
```

### Configuring Back-to-Back Frames Benchmarking Test Parameters on the ACX Series Router

---

#### Step-by-Step Procedure

The following configuration requires you to configure a test profile for the back to back frames test and reference the test-profile in a unique test-name. The test-name defines the parameters for the back to back frames test to be performed on the ACX router.

To configure the back-to-back frames test parameters on the ACX Router:

1. In configuration mode, at the **[edit]** hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```

2. Define a name for the back-to-back test profile—for example, b2bt.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit profiles test-profile b2bt
```

3. Configure the type of test to be performed as back-to-back frames, specify the packet size as 128 bytes, and define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles test-profile b2bt]
user@host# set test-type back-to-back-frames packet-size 4444 bandwidth-kbps
950000
```

4. Enter the **up** command twice to go to the **[edit services rpm rfc2544-benchmarking]** level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile b2bt]
user@host# up
user@host# up
```



5. Define a name for the back-to-back frames test—for example, b2bt-test. The test name can be up to 32 characters in length.  

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name b2bt-test
```
6. Specify the name of the test profile, b2bt, to be associated with the test name.  

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set test-profile b2bt
```
7. Configure the source and destination MAC address for the test packet.  

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set source-mac-address 00:00:00:00:11:11 destination-mac-address
00:00:00:00:22:22
```
8. Configure the outer VLAN ID for the test frames and specify the service type under test.  

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set ovlan-id 400 service-type elan
```
9. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.  

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set mode initiate-and-terminate
```
10. Configure the family type, **bridge**, for the benchmarking test and specify the direction, egress.  

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set family bridge direction egress
```
11. Specify the duration of each iteration in seconds, with a value from 10 seconds to 1,728,000 seconds. Also, specify the logical interface, ge-0/2/1.0, on which the RFC2544-based benchmarking test is run.  

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set test-iterator-duration 20 test-interface ge-1/1/3.0
```

### Configuring Latency Benchmarking Test Parameters on the ACX Series Router

#### Step-by-Step Procedure

The following configuration requires you to configure a test profile for the latency test and reference the test-profile in a unique test-name. The test-name defines the parameters for the latency test to be performed on the ACX router.

To configure the latency test parameters on the ACX Router:

1. In configuration mode, at the **[edit]** hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.  

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```
2. Define a name for the latency test profile—for example, lty.  

```
[edit services rpm rfc2544-benchmarking]
```

```
user@host# edit profiles test-profile lty
```

3. Configure the type of test to be performed as latency, specify the packet size of the test packet, and define the maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles]
```

```
user@host# set test-profile lty test-type latency packet-size 512 bandwidth-kbps 1000000
```

4. Enter the **up** command twice to go to the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile lty]
```

```
user@host# up
```

```
user@host# up
```

5. Define a name for the latency test—for example, lty-test. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
```

```
user@host# edit tests test-name lty-test
```

6. Specify the name of the test profile, lty, to be associated with the test name.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
```

```
user@host# set test-profile lty
```

7. Configure the source and destination MAC address for the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
```

```
user@host# set source-mac-address 00:00:00:00:11:11 destination-mac-address 00:00:00:00:22:22
```

8. Configure the outer VLAN ID for the test frames and specify the service type under test.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
```

```
user@host# set ovlan-id 400 service-type elan
```

9. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
```

```
user@host# set mode initiate-and-terminate
```

10. Configure the family type, **bridge**, for the benchmarking test and specify the direction, egress. Also, specify the source and destination UDP port to be used in the UDP headers of the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
```

```
user@host# set family bridge direction egress source-udp-port 200 destination-udp-port 200
```

11. Specify the duration of each iteration in seconds, with a value from 10 seconds to 1,728,000 seconds. Also, specify the logical interface, ge-0/2/1.0, on which the RFC2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
```

```
user@host# set test-iterator-duration 20 test-interface ge-1/1/3.0
```

### Configuring Frame Loss Benchmarking Test Parameters on the ACX Series Router

#### Step-by-Step Procedure

The following configuration requires you to configure a test profile for the frame loss test and reference the test-profile in a unique test-name. The test-name defines the parameters for the frame loss test to be performed on the ACX router.

To configure the frame loss test parameters on the ACX Router:

1. In configuration mode, at the `[edit]` hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```

2. Define a name for the frame loss test profile—for example, `frloss`.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit profiles test-profile frloss
```

3. Configure the type of test performed as frame loss, specify the packet size of the test packet, and define the maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles]
user@host# set test-profile frloss test-type frame-loss packet-size 1600
bandwidth-kbps 1000000
```

4. Enter the `up` command to go to the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles]
user@host# up
```

5. Define a name for the frame loss test—for example, `frloss-test`. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name frloss-test
```

6. Specify the name of the test profile, `frloss`, to be associated with the test name.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set test-profile frloss
```

7. Configure the source and destination MAC address for the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set source-mac-address 00:00:00:00:11:11 destination-mac-address
00:00:00:00:22:22
```

8. Configure the outer VLAN ID for the test frames and specify the service type under test.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set ovlan-id 400 service-type elan
```

9. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
```

```
user@host# set mode initiate-and-terminate
```

10. Configure the family type, **bridge**, for the benchmarking test and specify the direction, egress. Also, specify the source and destination UDP port to be used in the UDP headers of the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set family bridge direction egress source-udp-port 200
destination-udp-port 200
```

11. Specify the duration of each iteration in seconds, with a value from 10 seconds to 1,728,000 seconds. Also, specify the logical interface, ge-0/2/1.0, on which the RFC2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set test-iterator-duration 20 test-interface ge-1/1/3.0
```

12. Enter the **exit** command to go to the [edit] hierarchy level.

```
[edit services rpm rfc2544-benchmarking tests test-name test4]
user@host# exit
```

---

### Configuring Other Benchmarking Test Parameters on the ACX Series Router

---

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see Using the CLI Editor in Configuration Mode in the CLI User Guide.

To configure the interface and bridge domain on the ACX Router:

1. Configure the Layer 2 NNI interface on which the tests must be run from the **[edit]** hierarchy level.

```
[edit]
user@host# edit interfaces ge-1/2/1
```

2. Configure flexible VLAN tagging for the transmission of untagged frames or 802.1Q single-tagged and dual-tagged frames on the logical interface. You can also specify the maximum transmission unit (MTU) size for the interface and the encapsulation.

```
[edit interfaces ge-1/2/1]
user@host# set flexible-vlan-tagging mtu 9192 encapsulation
flexible-ethernet-services
```

3. Configure a logical unit for the interface, specify the encapsulation, and configure the VLAN ID on the logical interfaces.

```
[edit interfaces ge-1/2/1]
user@host# set unit 0 encapsulation vlan-bridge vlan-id 400
```

4. Configure the Layer 2 UNI interface.

```
[edit]
user@host# edit interfaces ge-1/1/3
```

5. Configure flexible VLAN tagging for transmission of non-tagged frames or 802.1Q single-tag and dual-tag frames on the logical interface. You can also specify the maximum transmission unit (MTU) size for the interface and the encapsulation.

```
[edit interfaces ge-1/1/3]
user@host# set flexible-vlan-tagging mtu 9192 encapsulation
flexible-ethernet-services
```

6. Configure a logical unit for the interface and specify the encapsulation and configure the VLAN ID on the logical interfaces.

```
[edit interfaces ge-1/1/3]
user@host# set unit 0 encapsulation vlan-bridge vlan-id 400
```

7. Configure the bridge domain, bd1, and specify the VLAN ID associated with the bridge domain and the associated interfaces from the [edit] hierarchy level.

```
[edit]
user@host# set bridge-domains bd1 vlan-id 600 interface ge-1/2/1.0
user@host# set bridge-domains bd1 vlan-id 600 interface ge-1/1/3.0
```

### Configuring Benchmarking Test Parameters on the MX104 Router

#### Step-by-Step Procedure

The following configuration requires you to configure a unique test-name for the benchmarking test on the MX104 router. The test-name defines the parameters for the benchmarking test to be performed. Because the test interface and test MAC addresses are the same, you can create a single test configuration at the reflector (MX104).

To configure the benchmarking test parameters on the MX104 Router:

1. In configuration mode, at the [edit] hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```

2. Define a name for the test—for example, l2b-reflector. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name l2b-reflector
```

3. Specify the source and destination MAC addresses of the test packet.

```
[edit services rpm rfc2544-benchmarking test-name l2b-reflector]
user@host# set source-mac-address 00:00:00:00:11:11 destination-mac-address
00:00:00:00:22:22
```

4. Specify the service type under test and the mode which is reflect, at the reflector.

```
[edit services rpm rfc2544-benchmarking test-name l2b-reflector]
user@host# set service-type elan
```

5. Specify the mode which is reflect at the reflector.

```
[edit services rpm rfc2544-benchmarking test-name l2b-reflector]
user@host# set mode reflect
```

6. Configure the family type, **bridge**, for the benchmarking test and specify the direction, egress. Also, specify the logical interface, ge-1/1/5.0, on which the RFC2544-based benchmarking test is being run.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-reflector]
```

```
user@host# set family bridge direction egress test-interface ge-1/1/5.0
```

### Configuring Other Benchmarking Test Parameters on the MX104 Router

#### **Step-by-Step Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the CLI User Guide.

To configure the interface and bridge domain on the MX104 Router:

1. Configure the Layer 2 NNI interface on which the tests must be run.  

```
[edit]
user@host# edit interfaces ge-1/1/6
```
2. Configure flexible VLAN tagging for transmission of untagged frames or 802.1Q single-tagged and dual-tagged frames on the logical interface. You can also specify the maximum transmission unit (MTU) size for the interface and the encapsulation.  

```
[edit interfaces ge-1/1/6]
user@host# set flexible-vlan-tagging mtu 9192 encapsulation
flexible-ethernet-services
```
3. Configure a logical unit for the interface, specify the encapsulation, and configure the VLAN ID on the logical interface.  

```
[edit interfaces ge-1/1/6]
user@host# set unit 0 encapsulation vlan-bridge vlan-id 400
```
4. Configure the Layer 2 NNI interface.  

```
[edit]
user@host# edit interfaces ge-1/1/5
```
5. Configure flexible VLAN tagging for transmission of untagged frames or 802.1Q single-tagged and dual-tagged frames on the logical interface. You can also specify the maximum transmission unit (MTU) size for the interface and the encapsulation.  

```
[edit interfaces ge-1/1/5]
user@host# set flexible-vlan-tagging mtu 9192 encapsulation
flexible-ethernet-services
```
6. Configure a logical unit for the interface, specify the encapsulation, and configure the VLAN ID on the logical interfaces.  

```
[edit interfaces ge-1/1/5]
user@host# set unit 0 encapsulation vlan-bridge vlan-id 400
```
7. Configure the bridge domain, bd1, and specify the VLAN ID associated with the bridge domain, and the associated interfaces from the [edit] hierarchy level.  

```
[edit]
user@host# set bridge-domains bd1 vlan-id 500 interface ge-1/1/6.0
user@host# set bridge-domains bd1 vlan-id 500 interface ge-1/1/5.0
```
8. Start the benchmarking test on the reflector.  

```
user@host> test services rpm rfc2544-benchmarking test l2b-reflector start
```

After the test is successfully completed at the initiator, you can stop the test at the reflector by entering the **test services rpm rfc2544-benchmarking test l2b-reflector stop** command.

## Results

---

In configuration mode, confirm your configuration on the ACX Router and the MX104 Router by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Benchmarking Test Parameters on the ACX Router :

```
[edit interfaces]
ge-1/2/1 {
 flexible-vlan-tagging;
 mtu 9192;
 encapsulation flexible-ethernet-services;
 unit 0 {
 encapsulation vlan-bridge;
 vlan-id 400;
 }
}
ge-1/1/3 {
 flexible-vlan-tagging;
 mtu 9192;
 encapsulation flexible-ethernet-services;
 unit 0 {
 encapsulation vlan-bridge;
 vlan-id 400;
 }
}

[edit bridge-domains]
bd1 {
 vlan-id 600;
 interface ge-1/2/1.0;
 interface ge-1/1/3.0;
}

[edit services rpm]
rfc2544-benchmarking {
 profiles {
 test-profile tput {
 test-type throughput
 packet-size 128;
 bandwidth-kbps 900000;
 }
 test-profile b2bt {
 test-type back-back-frames
 packet-size 512;
 bandwidth-kbps 950000;
 }
 test-profile lty {
 test-type latency
 }
 }
}
```

```
 packet-size 512;
 bandwidth-kbps 100000;
 }
test-profile frloss {
 test-type frameloss
 packet-size 1600;
 bandwidth-kbps 1000000;
}

tests {
 test-name tput-test {
 interface ge-1/1/3.0;
 test-profile tput;
 mode initiate,terminate;
 source-mac-address 00:00:00:00:11:11;
 destination-mac-address 00:00:00:00:22:22;
 ovlan-id 400;
 service-type elan;
 family bridge;
 direction egress;
 source-udp-port 200;
 destination-udp-port 200;
 test-iterator-duration 20;
 }
 test-name b2b-test {
 interface ge-1/1/3.0;
 test-profile b2bt;
 mode initiate,terminate;
 source-mac-address 00:00:00:00:11:11;
 destination-mac-address 00:00:00:00:22:22;
 ovlan-id 400;
 service-type elan;
 family bridge;
 direction egress;
 test-iterator-duration 20;
 }
 test-name lty-test {
 interface ge-1/1/3.0;
 test-profile lty;
 mode initiate,terminate;
 source-mac-address 00:00:00:00:11:11;
 destination-mac-address 00:00:00:00:22:22;
 ovlan-id 400;
 service-type elan;
 family bridge;
 direction egress;
 source-udp-port 200;
 destination-udp-port 200;
 test-iterator-duration 20;
 }
 test-name frloss-test {
 interface ge-1/1/3.0;
 test-profile frloss;
 mode initiate,terminate;
 source-mac-address 00:00:00:00:11:11;
```



```

 destination-mac-address 00:00:00:00:22:22;
 ovlan-id 400;
 service-type elan;
 family bridge;
 direction egress;
 source-udp-port 200;
 destination-udp-port 200;
 test-iterator-duration 20;
 }
}
}

```

Benchmarking Test Parameters on the MX104 Router:

```

[edit interfaces]
ge-1/1/6 {
 flexible-vlan-tagging;
 mtu 9192;
 encapsulation flexible-ethernet-services;
 unit 0 {
 encapsulation vlan-bridge;
 vlan-id 400;
 }
}
ge-1/1/5 {
 flexible-vlan-tagging;
 mtu 9192;
 encapsulation flexible-ethernet-services;
 unit 0 {
 encapsulation vlan-bridge;
 vlan-id 400;
 }
}
}
[edit bridge-domains]
bd1 {
 vlan-id 500;
 interface ge-1/1/6.0;
 interface ge-1/1/5.0;
}

[edit services rpm]
rfc2544-benchmarking {
 # Note, When in reflector mode, test profile is not needed
 tests {
 test-name l2b-reflector {
 interface ge-1/1/5.0;
 source-mac-address 00:00:00:00:11:11;
 destination-mac-address 00:00:00:00:22:22;
 }
 family bridge;
 mode reflect;
 service-type elan;
 family bridge;
 direction egress;
 }
}

```

```
 }
 }
}
```

## Verifying the Results of the Benchmarking Tests for Layer 2 Services (E-LAN) in Bridge Domains

Examine the results of the benchmarking tests that are performed on the configured service between the ACX Router and the MX104 Router. Start the test on the reflector first and then start the test on the initiator.

- [Verifying the Throughput Benchmarking Test Results on page 1032](#)
- [Verifying the Back-to-Back Benchmarking Test Results on page 1034](#)
- [Verifying the Frame Loss Benchmarking Test Results on page 1036](#)
- [Verifying the Latency Benchmarking Test Results on page 1038](#)

### Verifying the Throughput Benchmarking Test Results

**Purpose** Verify that the necessary and statistical values are displayed for the benchmarking tests that are run on the configured service between the ACX router and the MX104 router.

**Action** In operational mode, enter the **show services rpm rfc2544-benchmarking test-id test-id-number detail** command on the ACX router.

```
user@host> show services rpm rfc2544-benchmarking test-id 1 detail
Test information :
 Test id: 1, Test name: tput_test, Test type: Throughput
 Test mode: Initiate-and-Terminate
 Test packet size: 128
 Test state: TEST_STATE_COMPLETED
 Status: Test-Completed
 Test start time: 2014-09-24 22:21:09 PDT
 Test finish time: 2014-09-24 22:21:33 PDT
 Counters last cleared: Never

Test-profile Configuration:
 Test-profile name: tput
 Test packet size: 128
 Theoretical max bandwidth : 900000 kbps

Test Configuration:
 Test mode: Initiate-and-Terminate
 Duration in seconds: 20
 Test finish wait duration in seconds: 1
 Test family: Bridge
 Test iterator pass threshold: 0.50 %
 Test receive failure threshold: 0.00 %
 Test transmit failure threshold: 0.50 %

Bridge family Configuration:
 Interface : ge-1/1/3.0
 Test direction: Egress
 Source mac address: 00:00:00:00:11:11
 Destination mac address: 00:00:00:00:22:22
 Outer vlan-id: 400
 Outer vlan priority: 0
 Outer vlan cfi: 0
```

```

Outer tag protocol id: 0x8100
Source ipv4 address: 192.168.1.10
Destination ipv4 address: 192.168.1.20
Source udp port: 200
Destination udp port: 200

```

```

Rfc2544 throughput test information :
Initial test load percentage : 100.00 %
Test iteration mode : Binary
Test iteration step : 50.00 %
Theoretical max bandwidth : 900000 kbps

```

Test packet size: 128

Iteration	Internal Overhead	Duration (sec)	Elapsed time	Theoretical	Transmit	Measured
1	0	20	20	100.00 %	100.00 %	100.00 %

```

Result of the iteration runs : Throughput Test complete for packet size 128
Best iteration : 1, Best iteration (pps) : 760135
Best iteration throughput : 100.00 %

```

RFC2544 Throughput test results summary:

Packet Size	Internal overhead	Theoretical rate (pps)	Transmit pps	Tx Packets	Rx Packets	Measured throughput %
128	0	760135	760135	15202700	15202700	100.00 %
900000						

In operational mode, enter the **show services rpm rfc2544-benchmarking test-id test-id-number detail** command on the MX104 router.

```

user@host> show services rpm rfc2544-benchmarking test-id 1 detail
Test information :
Test id: 1, Test name: 12b-reflector, Test type: Reflect
Test mode: Reflect
Test packet size: 0
Test state: TEST_STATE_RUNNING
Status: Running
Test start time: 2014-09-24 22:20:54 PDT
Test finish time: TEST_RUNNING
Counters last cleared: Never

```

```

Test Configuration:
Test mode: Reflect
Duration in seconds: 864000
Test finish wait duration in seconds: 1
Test family: Bridge
Test iterator pass threshold: 0.50 %
Test receive failure threshold: 0.00 %
Test transmit failure threshold: 0.50 %

```

```

Bridge family Configuration:
Interface : ge-1/1/5.0
Test direction: Egress
Source mac address: 00:00:00:00:11:11
Destination mac address: 00:00:00:00:22:22
Service type: Elan

```

Elapsed time	Reflected Packets	Reflected Bytes
61	15202700	1945945600

You can also use the **show services rpm rfc2544-benchmarking (aborted-test | active-tests | completed-tests | summary)** command to display information about the results of each category or state of the RFC2544-based benchmarking tests for each real-time performance monitoring (RPM) instance.

**Meaning** The output displays the details of the benchmarking test that was performed. For more information about the **run show services rpm rfc2544-benchmarking** operational command, see **show services rpm rfc2544-benchmarking** in the CLI Explorer.

### Verifying the Back-to-Back Benchmarking Test Results

**Purpose** Verify that the necessary and statistical values are displayed for the benchmarking tests that are run on the configured service between the ACX router and the MX104 router.

**Action** In operational mode, enter the **show services rpm rfc2544-benchmarking test-id test-id-number detail** command on the ACX router.

```
user@host> show services rpm rfc2544-benchmarking test-id 4 detail
Test information :
 Test id: 4, Test name: b2b-test, Test type: Back-Back-Frames
 Test mode: Initiate-and-Terminate
 Test packet size: 128 512
 Test state: TEST_STATE_COMPLETED
 Status: Test-Completed
 Test start time: 2014-09-24 22:30:16 PDT
 Test finish time: 2014-09-24 22:31:03 PDT
 Counters last cleared: Never

Test-profile Configuration:
 Test-profile name: b2bt
 Test packet size: 128 512
 Theoretical max bandwidth : 950000 kbps

Test Configuration:
 Test mode: Initiate-and-Terminate
 Duration in seconds: 20
 Test finish wait duration in seconds: 1
 Test family: Bridge
 Test iterator pass threshold: 0.50 %
 Test receive failure threshold: 0.00 %
 Test transmit failure threshold: 0.50 %

Bridge family Configuration:
 Interface : ge-1/1/3.0
 Test direction: Egress
 Source mac address: 00:00:00:00:11:11
 Destination mac address: 00:00:00:00:22:22
 Outer vlan-id: 400
 Outer vlan priority: 0
 Outer vlan cfi: 0
 Outer tag protocol id: 0x8100
 Source ipv4 address: 192.168.1.10
 Destination ipv4 address: 192.168.1.20
```

Source udp port: 4040  
Destination udp port: 4041

Rfc2544 Back-Back test information :  
Initial burst length: 20 seconds at 950000 kbps  
Test iteration mode : Binary  
Test iteration step : 50.00 %

Test packet size: 128

Iteration	Theoretical burst length (packets)	Transmit burst length (packets)	Internal overhead	Duration time	Elapsed
1	16047280	16047280	0	20	20

Result of the iteration runs : Back-Back Test complete for packet size 128  
Best iteration : 1  
Measured burst (num sec) : 20 sec  
Measured burst (num pkts) : 16047280 packets

Test packet size: 512

Iteration	Theoretical burst length (packets)	Transmit burst length (packets)	Internal overhead	Duration time	Elapsed
1	4464280	4464280	0	20	20

Result of the iteration runs : Back-Back Test complete for packet size 512  
Best iteration : 1  
Measured burst (num sec) : 20 sec  
Measured burst (num pkts) : 4464280 packets

RFC2544 Back-Back test results summary:

Packet Size	Measured Burst length (Packets)	Time (seconds)
128	16047280	20
512	4464280	20

In operational mode, enter the **show services rpm rfc2544-benchmarking test-id test-id-number detail** command on the MX104 router.

```
user@host> show services rpm rfc2544-benchmarking test-id 4 detail
Test information :
 Test id: 4, Test name: l2b-reflector, Test type: Reflect
 Test mode: Reflect
 Test packet size: 0
 Test state: TEST_STATE_RUNNING
 Status: Running
 Test start time: 2014-09-24 22:30:07 PDT
 Test finish time: TEST_RUNNING
 Counters last cleared: Never

Test Configuration:
 Test mode: Reflect
 Duration in seconds: 864000
 Test finish wait duration in seconds: 1
 Test family: Bridge
 Test iterator pass threshold: 0.50 %
 Test receive failure threshold: 0.00 %
 Test transmit failure threshold: 0.50 %
```

```
Bridge family Configuration:
Interface : ge-1/1/5.0
Test direction: Egress
Source mac address: 00:00:00:00:11:11
Destination mac address: 00:00:00:00:22:22
Service type: Elan
```

Elapsed time	Reflected Packets	Reflected Bytes
58	20511560	4339763200

You can also use the **show services rpm rfc2544-benchmarking (aborted-test | active-tests | completed-tests | summary)** command to display information about the results of each category or state of the RFC2544-based benchmarking tests for each real-time performance monitoring (RPM) instance.

**Meaning** The output displays the details of the benchmarking test that was performed. For more information about the **run show services rpm rfc2544-benchmarking** operational command, see **show services rpm rfc2544-benchmarking** in the CLI Explorer.

---

### Verifying the Frame Loss Benchmarking Test Results

**Purpose** Verify that the necessary and statistical values are displayed for the benchmarking tests that are run on the configured service between the ACX router and the MX104 router.

**Action** In operational mode, enter the **show services rpm rfc2544-benchmarking test-id test-id-number detail** command on the ACX router.

```
user@host> show services rpm rfc2544-benchmarking test-id 3 detail
Test information :
 Test id: 3, Test name: frloss-test, Test type: Frame-Loss
 Test mode: Initiate-and-Terminate
 Test packet size: 1600
 Test state: TEST_STATE_COMPLETED
 Status: Test-Completed
 Test start time: 2014-09-24 22:26:45 PDT
 Test finish time: 2014-09-24 22:27:55 PDT
 Counters last cleared: Never
```

```
Test-profile Configuration:
 Test-profile name: frloss
 Test packet size: 1600
 Theoretical max bandwidth : 1000000 kbps
```

```
Test Configuration:
 Test mode: Initiate-and-Terminate
 Duration in seconds: 20
 Test finish wait duration in seconds: 1
 Test family: Bridge
 Test iterator pass threshold: 0.50 %
 Test receive failure threshold: 0.00 %
 Test transmit failure threshold: 0.50 %
```

```
Bridge family Configuration:
Interface : ge-1/1/3.0
Test direction: Egress
Source mac address: 00:00:00:00:11:11
```

```

Destination mac address: 00:00:00:00:22:22
Outer vlan-id: 400
Outer vlan priority: 0
Outer vlan cfi: 0
Outer tag protocol id: 0x8100
Source ipv4 address: 192.168.1.10
Destination ipv4 address: 192.168.1.20
Source udp port: 200
Destination udp port: 200

```

```

Rfc2544 frame-loss test information :
Initial test load percentage : 100.00 %
Test iteration mode : step-down
Test iteration step : 10 %
Theoretical max bandwidth : 1000000 kbps

```

Test packet size: 1600

Iteration	Internal Duration	Elapsed	-----	Throughput	-----	Frame-loss
	Overhead (sec)	time		Theoretical	Transmit Measured	rate %
1	0	20	20	100.00 %	100.00 % 100.00 %	0.00 %
2	0	20	20	100.00 %	100.00 % 100.00 %	0.00 %
3	0	20	20	100.00 %	100.00 % 100.00 %	0.00 %

```

Result of the iteration runs : Frame-loss test complete for packet size 1600
Percentage throughput transmitted: 100.00 %
Frame-loss rate (percent) : 0.00 %

```

RFC2544 Frame-loss test results summary:

Packet Size	Internal Frame Loss overhead rate percent	Theoretical rate (pps)	Transmit pps	Transmit throughput	Tx Packets	Rx Packets
1600	0	77160	77160	100.00 %	1543200	1543200
	0.00 %					

In operational mode, enter the **show services rpm rfc2544-benchmarking test-id test-id-number detail** command on the MX104 router.

```

user@host> show services rpm rfc2544-benchmarking test-id 3 detail
Test information :
 Test id: 3, Test name: l2b-reflector, Test type: Reflect
 Test mode: Reflect
 Test packet size: 0
 Test state: TEST_STATE_RUNNING
 Status: Running
 Test start time: 2014-09-24 22:25:36 PDT
 Test finish time: TEST_RUNNING
 Counters last cleared: Never

Test Configuration:
 Test mode: Reflect
 Duration in seconds: 864000
 Test finish wait duration in seconds: 1
 Test family: Bridge
 Test iterator pass threshold: 0.50 %
 Test receive failure threshold: 0.00 %
 Test transmit failure threshold: 0.50 %

```

```
Bridge family Configuration:
Interface : ge-1/1/5.0
Test direction: Egress
Source mac address: 00:00:00:00:11:11
Destination mac address: 00:00:00:00:22:22
Service type: Elan
```

Elapsed time	Reflected Packets	Reflected Bytes
95	1624361	2598977600

You can also use the **show services rpm rfc2544-benchmarking (aborted-test | active-tests | completed-tests | summary)** command to display information about the results of each category or state of the RFC2544-based benchmarking tests for each real-time performance monitoring (RPM) instance.

**Meaning** The output displays the details of the benchmarking test that was performed. For more information about the **run show services rpm rfc2544-benchmarking** operational command, see **show services rpm rfc2544-benchmarking** in the CLI Explorer.

---

### Verifying the Latency Benchmarking Test Results

**Purpose** Verify that the necessary and statistical values are displayed for the benchmarking tests that are run on the configured service between the ACX router and the MX104 router.

**Action** In operational mode, enter the **show services rpm rfc2544-benchmarking test-id test-id-number detail** command on the ACX router.

```
user@host> show services rpm rfc2544-benchmarking test-id 5 detail
Test information :
```

```
Test id: 5, Test name: lty-test, Test type: Latency
Test mode: Initiate-and-Terminate
Test packet size: 512
Test state: TEST_STATE_COMPLETED
Status: Test-Completed
Test start time: 2014-09-24 22:33:05 PDT
Test finish time: 2014-09-24 22:40:46 PDT
Counters last cleared: Never
```

```
Test-profile Configuration:
Test-profile name: lty
Test packet size: 512
Theoretical max bandwidth : 1000000 kbps
```

```
Test Configuration:
Test mode: Initiate-and-Terminate
Duration in seconds: 20
Test finish wait duration in seconds: 1
Test family: Bridge
Test iterator pass threshold: 0.50 %
Test receive failure threshold: 0.00 %
Test transmit failure threshold: 0.50 %
```

```
Bridge family Configuration:
Interface : ge-1/1/3.0
Test direction: Egress
Source mac address: 00:00:00:00:11:11
```



```

Destination mac address: 00:00:00:00:22:22
Outer vlan-id: 400
Outer vlan priority: 0
Outer vlan cfi: 0
Outer tag protocol id: 0x8100
Source ipv4 address: 192.168.1.10
Destination ipv4 address: 192.168.1.20
Source udp port: 200
Destination udp port: 200

```

```

Rfc2544 latency test information :
Theoretical max bandwidth : 1000000 kbps
Initial test load percentage : 100.00 %
Duration in seconds: 20
Measurement unit for timestamp: Nanoseconds

```

Test packet size: 512

Iteration	Duration	Elapsed	Theoretical	Transmit	Throughput	
		Latency				
	(sec)	time	rate (pps)	pps	percent	Minimum
	Average	Maximum	Probe			
1	20	20	234962	234962	100.00 %	44008
	45253	47424	45096			
2	20	20	234962	234962	100.00 %	44008
	45237	47456	45256			
3	20	20	234962	234962	100.00 %	43864
	45198	46976	45144			
4	20	20	234962	234962	100.00 %	43832
	45243	47088	45096			
5	20	20	234962	234962	100.00 %	44072
	45261	46976	45176			
6	20	20	234962	234962	100.00 %	43784
	45214	46864	45032			
7	20	20	234962	234962	100.00 %	44024
	45259	47216	45240			
8	20	20	234962	234962	100.00 %	44072
	45290	46864	45192			
9	20	20	234962	234962	100.00 %	43976
	45272	46792	45208			
10	20	20	234962	234962	100.00 %	44024
	45206	46976	45112			
11	20	20	234962	234962	100.00 %	44040
	45198	47088	45176			
12	20	20	234962	234962	100.00 %	44008
	45223	46976	45160			
13	20	20	234962	234962	100.00 %	44088
	45257	47408	45176			
14	20	20	234962	234962	100.00 %	43976
	45183	46832	45080			
15	20	20	234962	234962	100.00 %	44024
	45198	47088	45112			
16	20	20	234962	234962	100.00 %	43864
	45206	46912	45208			
17	20	20	234962	234962	100.00 %	44056
	45209	46960	45176			
18	20	20	234962	234962	100.00 %	44008
	45198	46912	45112			
19	20	20	234962	234962	100.00 %	43816
	45175	47248	45000			
20	20	20	234962	234962	100.00 %	43912
	45202	46992	45192			

Result of the iteration runs : Latency Test complete for packet size 512  
 Internal overhead per packet: 0  
 Avg (min) Latency : 43972  
 Avg (avg) latency : 45224  
 Avg (Max) latency : 47052  
 Avg (probe) latency : 45147

RFC2544 Latency test results summary:

```

Packet Internal Theoretical Transmit Tx Rx
----- Latency -----
Size overhead rate (pps) pps Packets Packets Minimum
Average Maximum Probe
512 0 234962 234962 93984800 93984800 43972
45224 47052 45147
```

In operational mode, enter the **show services rpm rfc2544-benchmarking test-id test-id-number detail** command on the MX104 router.

```
user@host> show services rpm rfc2544-benchmarking test-id 5 detail
Test information :
 Test id: 5, Test name: 12b-reflector, Test type: Reflect
 Test mode: Reflect
 Test packet size: 0
 Test state: TEST_STATE_RUNNING
 Status: Running
 Test start time: 2014-09-24 22:32:55 PDT
 Test finish time: TEST_RUNNING
 Counters last cleared: Never
```

```
Test Configuration:
 Test mode: Reflect
 Duration in seconds: 864000
 Test finish wait duration in seconds: 1
 Test family: Bridge
 Test iterator pass threshold: 0.50 %
 Test receive failure threshold: 0.00 %
 Test transmit failure threshold: 0.50 %
```

```
Bridge family Configuration:
 Interface : ge-1/1/5.0
 Test direction: Egress
 Source mac address: 00:00:00:00:11:11
 Destination mac address: 00:00:00:00:22:22
 Service type: Elan
```

```
Elapsed Reflected Reflected
time Packets Bytes
426 84586320 43308195840
```

You can also use the **show services rpm rfc2544-benchmarking (aborted-test | active-tests | completed-tests | summary)** command to display information about the results of each category or state of the RFC2544-based benchmarking tests for each real-time performance monitoring (RPM) instance.

**Meaning** The output displays the details of the benchmarking test that was performed. For more information about the **run show services rpm rfc2544-benchmarking** operational command, see **show services rpm rfc2544-benchmarking** in the CLI Explorer.

**Related Documentation**

- [Layer 2 RFC2544-Based Benchmarking Tests Overview on page 986](#)
- [Supported RFC2544-Based Benchmarking Statements on MX104 Routers on page 988](#)



# Tracking Streaming Media Traffic Using Inline Video Monitoring

- [Inline Video Monitoring Overview on page 1043](#)
- [Configuring Inline Video Monitoring on page 1045](#)
- [Inline Video Monitoring Syslog Messages on page 1047](#)

## Inline Video Monitoring Overview

---

Junos OS supports inline video monitoring using Media Delivery Index (MDI) metrics.

Inline video monitoring is available on MX Series routers using only the following MPCs:

- MPCE1
- MPCE2
- MPC-16XGE

You use the **video-monitoring** statement at the **[edit services]** hierarchy level to specify monitoring criteria for two key indicators of video traffic problems: delay factor and media loss rate (MLR), and to apply these metrics to flows on designated interfaces.

Before you use the inline video monitoring feature, ensure that you understand the following terms:

- **media delivery index**—These metrics facilitate identification of buffering needs for streaming media. Buffering must be adequate to compensate for packet jitter, measured by the MDI delay factor, and quality problems indicated by lost packets, measured by the MDI media loss rate (MLR). By performing measurements under varying load conditions, you can identify sources of significant jitter or packet loss and take appropriate action.
- **delay factor** —Delay factor is the maximum observed time difference between the arrival of media data and the drain of media data. The expected drain rate is the nominal, constant traffic rate for constant bit rate streams or the computed traffic rate of variable rate media stream packet data.

For typical stream rates of 1 megabit per second and higher, an interval of one second provides an adequate sample time. The delay factor indicates how long a data stream must be buffered (delayed) at its nominal bit rate to prevent packet loss.

The delay factor suggests the minimum size of the buffer required at the next downstream node. As a stream progresses, the variation of the delay factor indicates packet bunching or packet gaps (jitter). Greater delay factor values also indicate that more network latency is needed to deliver a stream due to the need to pre-fill a receive buffer before beginning the drain to guarantee no underflow.

When the nominal drain bit rate at a receiving node is known, the delay factor's maximum indicates the size of buffer required to accommodate packet jitter.

- **Media rate variation (MRV)**—This value is the difference between the expected packet rate and actual packet rate expressed as a percentage of the expected packet rate.
- **Media loss rate (MLR)**—This value is the number of media packets lost over a configurable time interval (*interval-duration*,) where the flow packets are packets carrying streaming application information. A single IP packet can contain zero or more streaming packets. For example, an IP packet typically contains seven 188-byte MPEG transport stream packets. In this case, a single IP packet loss results in seven lost packets counted (if those seven lost packets did not include null packets). Including out-of-order packets is important, because many stream consumer-type devices do not attempt to reorder packets that are received out of order.

To configure the monitoring process, define criteria templates and apply them to the interfaces and flows you want to monitor. Monitoring templates include the following criteria:

- Duration of each measurement cycle
- Flow rate information used to establish expected flow rates
- Threshold levels for media rate variation and media loss rate that trigger desired syslog alerts

For each interface you want to monitor, you can define one or more filters to select flows for monitoring. Flows are designated as input or output flows and are uniquely identified by:

- Source IP address
- Source port
- Destination IP address
- Destination port

Junos OS supports the definition of filters for up to 256 flows, which can consist of input flows, output flows, or a combination of input and output flows. These filters provide criteria for selecting flows for monitoring. If the selection criteria consist of lists of IP addresses or ports, you could exceed the maximum number of match conditions for flows. Video monitoring selects a widely variable number of flows based on flow filters. The total number of flows that can be measured at a time depends on the specific MPC card being used, as shown in [Table 41 on page 1045](#).

When you do not define input or output flow filters for a monitored interfaces, all flows on the interface are subject to monitoring.

**Table 41: MPC Flow Monitoring Capacity by Model**

MPC Model	Maximum Number of Flows Monitored Simultaneously
MPCE1	1000
MPCE2	2000
MPC-16XGE	4000



**NOTE:** Junos OS measures both UDP flows (the default) and RTP flows. Junos OS differentiates media traffic over UDP or RTP by inspecting the first byte in the UDP payload. If the first byte of the UDP payload is 0x47 (MPEG2-TS sync byte), the traffic is treated as media traffic over UDP. Traffic is treated as media traffic over RTP if the version field is 2 and the payload type is 33 in the RTP header. When neither of these criteria are met, the packet is not considered for video monitoring.

**Related Documentation**

- [Configuring Inline Video Monitoring on page 1045](#)
- [show services video-monitoring mdi stats fpc-slot on page 2291](#)
- [show services video-monitoring mdi errors fpc-slot on page 2285](#)
- [show services video-monitoring mdi flows fpc-slot on page 2287](#)

## Configuring Inline Video Monitoring

To configure inline video monitoring, perform the following tasks.

- [Configuring Media Delivery Indexing Criteria on page 1045](#)
- [Configuring Interface Flow Criteria on page 1047](#)

### Configuring Media Delivery Indexing Criteria

To configure media delivery indexing criteria:

1. In edit mode, create a named template for video monitoring.  

```
user@host# edit services video-monitoring templates template-name
```

For example,

```
user@host# edit services video-monitoring templates t1
```

2. Set the duration for sampling in seconds. Flow media delivery indexing statistics are updated at the end of this interval.

```
[edit services video-monitoring templates t1]
```

```
user@host# set interval-duration 1
```



**BEST PRACTICE:** If you change the interval duration when a template is being used, you cause a change in the calculated number of expected packets in an measurement interval for the template. We recommend that you do not change the interval duration for a template that is in use.

3. Set the inactivity timeout.

```
[edit services video-monitoring templates t1]
user@host# set inactivity-timeout 30
```

4. Configure either **media-rate** or **layer3-packet-rate** to establish expected flow rates used to compare to monitored flow rates.



**NOTE:** The media rate is the configured media bit rate for the stream. The media rate is used to establish *expected packets per second (pps)*.

The layer 3 packet rate in packets per second (pps) and is used to establish *expected bits per second (bps)*.

```
[edit services video-monitoring templates t1]
user@host# set media-rate 2972400
```

5. Set delay factor thresholds for syslog message levels.

```
[edit services video-monitoring templates t1]
user@host# set delay-factor threshold info 100
user@host# set delay-factor threshold warning 200
user@host# set delay-factor threshold critical 300
```

6. Set media loss rate thresholds for syslog message levels. You can set the threshold based on number of packets lost, or percentage of packets lost.

Or

```
[edit services video-monitoring templates t1]
user@host# set media-loss-rate threshold info percentage 5
user@host# set media-loss-rate threshold warning percentage 10
user@host# set media-loss-rate threshold critical percentage 20
```

7. Set the media rate variation thresholds for syslog message levels. The threshold is based on the ratio of the *difference* between the configured media rate and the monitored media rate to the configured media rate, expressed as a percentage.

```
[edit services video-monitoring templates t1]
user@host# set media-rate-variation threshold info 10
user@host# set media-rate-variation threshold warning 15
user@host# set media-rate-variation threshold critical 20
```



## Configuring Interface Flow Criteria

To configure monitoring of flows for interfaces:

1. In edit mode, identify an interface for monitoring .

```
user@host# edit services video-monitoring interfaces interface-name
```

2. Identify input flows for monitoring. Flows are uniquely identified by source IP address, source port, destination IP address, and destination port. You can restrict flow measurement by specifying values for these identifiers. You can specify individual addresses or ports or lists of addresses and ports. If you do not specify any identifiers, all flows on the interface are monitored.

```
[edit services video-monitoring interfaces interface-name]
user@host# set input-flows input-flow-name
user@host# set input-flows input-flow-name source-address address
user@host# set input-flows input-flow-name source-port port
user@host# set input-flows input-flow-name destination-address address
user@host# set input-flows input-flow-name destination-port port
```



**NOTE:** You can configure a maximum of 256 flow definitions. If your flow definitions contain lists of addresses and ports, you may exceed the number of match conditions. When you exceed the limits for flows or match conditions, you receive the following constraint message when you commit:

```
'interfaces xe-0/2/2.0'
 Number of flows or Number of match condition under flows exceeded
 limit
error: configuration check-out failed
```

3. Identify output flows for monitoring, using the same options listed in Step 2.
4. Identify the template used to monitor the flows on the interface.

```
[edit services video-monitoring interfaces interface-name]
set template t1
```

### Related Documentation

- [Inline Video Monitoring Overview on page 1043](#)
- [templates on page 1760](#)
- [interfaces on page 1694](#)

## Inline Video Monitoring Syslog Messages

The following examples show the syslog messages produced when configured video monitoring thresholds are exceeded.

`/var/log/messages`

```
Mar 11 18:36:25 tstrtr01 fpc2 [MDI] DF: 56.71 ms, exceeded threshold for
flow(src:20.0.0.2 dst:30.0.0.2 sport:1024 dport:2048) ingressing at interface
```

```
xe-2/2/1.0 with template t1.
Mar 11 18:36:25 tstrtr01 fpc2 [MDI] MLR : 112, exceeded threshold for flow
(src:20.0.0.2 dst:30.0.0.2 sport:1024 dport:2048) ingressing at interface
xe-2/2/1.0 with template t1.
Mar 11 18:36:25 tstrtr01 fpc2 [MDI] MRV : -5.67, exceeded threshold for flow
(src:20.0.0.2 dst:30.0.0.2 sport:1024 dport:2048) ingressing at interface
xe-2/2/1.0 with template t1.
```

### Console Messages

```
NPC2(tstrtr01 vty)# [Mar 12 01:40:58.411 LOG: Critical] [MDI] MLR : 420, exceeded
threshold for flow (src:20.0.0.2 dst:30.0.0.2 sport:1024 dport:2048) ingressing
at interface xe-2/2/1.0 with template t1.
[Mar 12 01:40:58.411 LOG: Critical] [MDI] MRV : -14.89, exceeded threshold for
flow (src:20.0.0.2 dst:30.0.0.2 sport:1024 dport:2048) ingressing at interface
xe-2/2/1.0 with template t1.
[Mar 12 01:40:59.412 LOG: Critical] [MDI] DF: 141.74 ms, exceeded threshold for
flow(src:20.0.0.2 dst:30.0.0.2 sport:1024 dport:2048) ingressing at interface
xe-2/2/1.0 with template t1.
```

**Related Documentation**

- [Configuring Inline Video Monitoring on page 1045](#)

## PART 18

# Sampling, Discard Accounting, and Port Mirroring Services

- [Sampling Data Using Traffic Sampling and Discard Accounting on page 1051](#)
- [Sampling Data Using Inline Sampling on page 1065](#)
- [Sampling Data Using Flow Aggregation on page 1077](#)
- [Sending Packets for Analysis Using Port Mirroring on page 1111](#)



# Sampling Data Using Traffic Sampling and Discard Accounting

- [Configuring Traffic Sampling on page 1051](#)
- [Sampling Instance Configuration on page 1061](#)
- [Configuring Discard Accounting on page 1063](#)

## Configuring Traffic Sampling

---

Traffic sampling enables you to copy traffic to a Physical Interface Card (PIC) that performs flow accounting while the router forwards the packet to its original destination. You can configure the router to perform sampling in either of two locations:

- On the Routing Engine, using the sampled process. To select this method, use a filter (input or output) with a matching term that contains the **then sample** statement.
- On the Monitoring Services, Adaptive Services, or Multiservices PIC.



**NOTE:** Routing Engine based sampling is not supported on VPN routing and forwarding (VRF) instances.

The following sections provide configuration instructions for traffic sampling:

- [Configuring Firewall Filter for Traffic Sampling on page 1051](#)
- [Configuring Traffic Sampling on a Logical Interface on page 1053](#)
- [Disabling Traffic Sampling on page 1054](#)
- [Sampling Once on page 1054](#)
- [Preserving Prerewrite ToS Value for Egress Sampled or Mirrored Packets on page 1055](#)
- [Configuring Traffic Sampling Output on page 1056](#)
- [Tracing Traffic Sampling Operations on page 1058](#)
- [Traffic Sampling Examples on page 1058](#)

## Configuring Firewall Filter for Traffic Sampling

To configure firewall filter for traffic sampling, you must perform the following tasks:

- Create a firewall filter to apply to the logical interfaces being sampled by including the **filter** statement at the **[edit firewall family *family-name*]** hierarchy level. In the filter **then** statement, you must specify the action modifier **sample** and the action **accept**.

```
filter filter-name {
 term term-name {
 then {
 sample;
 accept;
 }
 }
}
```

For more information about firewall filter actions and action modifiers, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

- Apply the filter to the interfaces on which you want to sample traffic by including the **address** and **filter** statements at the **[edit interfaces *interface-name* unit *logical-unit-number* family *family-name*]** hierarchy level:

```
address address {
}
filter {
 input filter-name;
}
```

The following prerequisites apply to M, MX, and T Series routers when you configure traffic sampling on interfaces and in firewall filters:

- If you configure a sample action in a firewall filter for an inet or inet6 family on an interface without configuring the forwarding-options settings, operational problems might occur if you also configure port mirroring or flow-tap functionalities. In such a scenario, all the packets that match the firewall filter are incorrectly sent to the service PIC.
- If you include the **then sample** statement at the **[edit firewall family inet filter *filter-name* term *term-name*]** hierarchy level to specify a sample action in a firewall filter for IPv4 packets, you must also include the **family inet** statement at the **[edit forwarding-options sampling]** hierarchy level or the **instance *instance-name* family inet** statement at the **[edit forwarding-options sampling]** hierarchy level. Similarly, if you include the **then sample** statement at the **[edit firewall family inet6 filter *filter-name* term *term-name*]** hierarchy level to specify a sample action in a firewall filter for IPv6 packets, you must also include **family inet6** statement at the **[edit forwarding-options sampling]** hierarchy level or the **instance *instance-name* family inet6** statement at the **[edit forwarding-options sampling]** hierarchy level. Otherwise, a commit error occurs when you attempt to commit the configuration.
- Also, if you configure traffic sampling on a logical interface by including the sampling input or sampling output statements at the **[edit interface *interface-name* unit *logical-unit-number*]** hierarchy level, you must also include the **family inet | inet6** statement at the **[edit forwarding-options sampling]** hierarchy level, or the **instance *instance-name* family inet | inet6** statement at the **[edit forwarding-options sampling]** hierarchy level.

## Configuring Traffic Sampling on a Logical Interface

To configure traffic sampling on any logical interface, enable sampling and specify a non zero sampling rate by including the sampling statement at the **[edit forwarding-options]** hierarchy level:

```
sampling {
 input {
 rate number;
 run-length number;
 max-packets-per-second number;
 maximum-packet-length bytes;
 }
}
```

When you use Routing Engine-based sampling, specify the threshold traffic value by including the **max-packets-per-second** statement. The value is the maximum number of packets to be sampled, beyond which the sampling mechanism begins dropping packets. The range is from 0 through 65,535. A value of 0 instructs the Packet Forwarding Engine not to sample any packets. The default value is 1000.




**NOTE:** When you configure active monitoring and specify a Monitoring Services, Adaptive Services, or Multiservices PIC in the output statement, the **max-packets-per-second** value is ignored.

Specify the sampling rate by setting the values for **rate** and **run-length** (see [Figure 34 on page 873](#)).

**Figure 41: Configuring Sampling Rate**

### Rate and Run-length

Case #1 Rate = 10, run-length = 4



1746



**NOTE:** If PIC-based flow monitoring is enabled on an *ms-fpc/pic/port.logical-unit* interface, a commit check error occurs when you attempt to configure ingress traffic sampling on that interface. This error occurs because a combination of ingress sampling and PIC-based flow monitoring operations on an *ms-* logical interface causes undesired flow monitoring behavior and might result in repeated sampling of a single packet. You must not configure ingress sampling on *ms-* logical interfaces on which PIC-based flow monitoring is enabled.

The **rate** statement specifies the ratio of packets to be sampled. For example, if you configure a rate of 10, *x* number of packets out of every 10 is sampled, where *x*=run length + 1. By default, the rate is 0, which means that no traffic is sampled.

The **run-length** statement specifies the number of matching packets to sample following the initial one-packet trigger event. By default, the run length is 0, which means that no more traffic is sampled after the trigger event. The range is from 0 through 20. Configuring a run length greater than 0 allows you to sample packets following those already being sampled.



**NOTE:** The **run-length** and **maximum-packet-length** configuration statements are not supported on MX80 routers.

If you do not include the **input** statement, sampling is disabled.

To collect the sampled packets in a file, include the **file** statement at the **[edit forwarding-options sampling output]** hierarchy level. Output file formats are discussed later in the chapter.

## Disabling Traffic Sampling

To explicitly disable traffic sampling on the router, include the **disable** statement at the **[edit forwarding-options sampling]** hierarchy level:

```
disable;
```

## Sampling Once

To explicitly sample a packet for active monitoring only once, include the **sample-once** statement at the **[edit forwarding-options sampling]** hierarchy level:

```
sample-once;
```

Setting this option avoids duplication of packets in cases where sampling is enabled at both the ingress and egress interfaces and simplifies analysis of the sampled traffic.



## Preserving Prerewrite ToS Value for Egress Sampled or Mirrored Packets

To preserve the prenormalized type-of-service (ToS) value in egress sampled or mirrored packets, include the **pre-rewrite-tos** statement at the **[edit forwarding-options sampling]** hierarchy level.

On MPC-based interfaces, you can configure ToS rewrite either using class-of-service (CoS) configuration by including the **rewrite-rules dscp rule\_name** statement at the **[edit class-of-service interfaces interface-name unit logical-unit-number]** hierarchy level or using firewall filter configuration by including the **dscp** statement at the **[edit firewall family family-name filter filter-name term term-name then]** hierarchy level. If ToS rewrite is configured, the egress mirrored or sampled copies contain the post-rewrite ToS values by default. With the **pre-rewrite-tos** configuration, you can retain the prerewrite ToS value in the sampled or mirrored packets.



### NOTE:

- If ToS rewrite is configured on the egress interface by using both CoS and firewall filter configuration, and if the **pre-rewrite-tos** statement is also configured, then the egress sampled packets contain the DSCP value set using the firewall filter configuration. However, if the **pre-rewrite-tos** statement is not configured, the egress sampled packets contain the DSCP value set by the CoS configuration.
- With the **pre-rewrite-tos** statement, you can configure retaining prenormalization ToS values only for sampling done under family inet and family inet6.
- This feature cannot be configured at the **[edit logical-systems]** hierarchy level. It can be configured only at the global level under the forwarding-option configuration.
- When ToS rewrite is configured by using a firewall filter on both ingress and egress interfaces, the egress sampled packets contain the DSCP value set by the ingress ToS rewrite configuration if the **pre-rewrite-tos** statement is configured. However, if the **pre-rewrite-tos** statement is not configured, the egress sampled packets contain the DSCP value set by the ToS rewrite configuration for the egress firewall filter.
- If the **pre-rewrite-tos** statement is configured, and a deactivate or delete operation is performed at the **[edit forwarding-options]** hierarchy level, **pre-rewrite-tos** configuration still remains active. To disable the **pre-rewrite-tos** configuration for such a case, you must explicitly deactivate or delete the **pre-rewrite-tos** statement at the **[edit forwarding-options sampling]** hierarchy level before performing a deactivate or delete operation at the **[edit forwarding-options]** hierarchy level.

## Configuring Traffic Sampling Output

To configure traffic sampling output, include the following statements at the **[edit forwarding-options sampling family (inet | inet6 | mpls) output]** hierarchy level:

```
aggregate-export-interval seconds;
flow-active-timeout seconds;
flow-inactive-timeout seconds;
extension-service service-name;
flow-server hostname {
 aggregation {
 autonomous-system;
 destination-prefix;
 protocol-port;
 source-destination-prefix {
 caida-compliant;
 }
 source-prefix;
 }
 autonomous-system-type (origin | peer);
 (local-dump | no-local-dump);
 port port-number;
 source-address address;
 version format;
 version9 {
 template template-name;
 }
}
interface interface-name {
 engine-id number;
 engine-type number;
 source-address address;
}
file {
 disable;
 filename filename;
 files number;
 size bytes;
 (stamp | no-stamp);
 (world-readable | no-world-readable);
}
```

To configure inline flow monitoring on MX Series routers, include the **inline-jflow** statement at the **[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output]** hierarchy level. Inline sampling exclusively supports a new format called IP\_FIX that uses UDP as the transport protocol. When you configure inline sampling, you must include the **version-ipfix** statement at the **[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output flow-server address]** hierarchy level and also at the **[edit services flow-monitoring]** hierarchy level. For more information about configuring inline flow monitoring, see [“Configuring Inline Active flow Monitoring” on page 890](#).

To direct sampled traffic to a flow-monitoring interface, include the **interface** statement. The **engine-id** and **engine-type** statements specify the identity and type numbers of the

interface; they are dynamically generated based on the Flexible PIC Concentrator (FPC), PIC, and slot numbers and the chassis type. The **source-address** statement specifies the traffic source.

To configure flow sampling version 9 output, you need to include the **template** statement at the **[edit forwarding-options sampling output version9]** hierarchy level. For information on cflowd, see [“Enabling Flow Aggregation” on page 898](#).

The **aggregate-export-interval** statement is described in [“Configuring Discard Accounting” on page 883](#), and the **flow-active-timeout** and **flow-inactive-timeout** statements are described in [“Configuring Flow Monitoring” on page 818](#).

Traffic sampling results are automatically saved to a file in the **/var/tmp** directory. To collect the sampled packets in a file, include the **file** statement at the **[edit forwarding-options sampling family inet output]** hierarchy level:

```
file {
 disable;
 filename filename;
 files number;
 size bytes;
 (stamp | no-stamp);
 (world-readable | no-world-readable);
}
```

### Traffic Sampling Output Format

Traffic sampling output is saved to an ASCII text file. The following is an example of the traffic sampling output that is saved to a file in the **/var/tmp** directory. Each line in the output file contains information for one sampled packet. You can optionally display a timestamp for each line.

The column headers are repeated after each group of 1000 packets.

```
Apr 7 15:48:50
Time Dest Src Dest Src Proto TOS Pkt Intf IP TCP
 addr addr port port len num frag flags
Apr 7 15:48:54 192.168.9.194 192.168.9.195 0 0 1 0x0 84 8 0x0 0x0
Apr 7 15:48:55 192.168.9.194 192.168.9.195 0 0 1 0x0 84 8 0x0 0x0
Apr 7 15:48:56 192.168.9.194 192.168.9.195 0 0 1 0x0 84 8 0x0 0x0
Apr 7 15:48:57 192.168.9.194 192.168.9.195 0 0 1 0x0 84 8 0x0 0x0
Apr 7 15:48:58 192.168.9.194 192.168.9.195 0 0 1 0x0 84 8 0x0 0x0
```

To set the timestamp option for the file **my-sample**, enter the following:

```
[edit forwarding-options sampling output file]
user@host# set filename my-sample files 5 size 2m world-readable stamp;
```

Whenever you toggle the timestamp option, a new header is included in the file. If you set the **stamp** option, the **Time** field is displayed.

```
Apr 7 15:48:50
Time Dest Src Dest Src Proto TOS Pkt Intf IP TCP
addr addr port port len num frag flags
Feb 1 20:31:21
```

#	Dest	Src	Dest	Src	Proto	TOS	Pkt	Intf	IP	TCP
#	addr	addr	port	port			len	num	frag	flags

## Tracing Traffic Sampling Operations

Tracing operations track all traffic sampling operations and record them in a log file in the `/var/log` directory. By default, this file is named `/var/log/sampled`. The default file size is 128K, and 10 files are created before the first one gets overwritten.

To trace traffic sampling operations, include the **traceoptions** statement at the **[edit forwarding-options sampling]** hierarchy level:

```
traceoptions {
 no-remote-trace;
 file filename <files number> <size bytes> <match expression> <world-readable |
 no-world-readable>;
}
```

## Traffic Sampling Examples

The following sections provide examples of configuring traffic sampling:

- [Example: Sampling a Single SONET/SDH Interface on page 1058](#)
- [Example: Sampling All Traffic from a Single IP Address on page 1059](#)
- [Example: Sampling All FTP Traffic on page 1060](#)

---

### Example: Sampling a Single SONET/SDH Interface

The following configuration gathers statistical sampling information from a small percentage of all traffic on a single SONET/SDH interface and collects it in a file named **sonet-samples.txt**.

Create the filter:

```
[edit firewall family inet]
filter {
 input sample-sonet {
 then {
 sample;
 accept;
 }
 }
}
```

Apply the filter to the SONET/SDH interface:

```
[edit interfaces]
so-0/0/1 {
 unit 0 {
 family inet {
 filter {
 input sample-sonet;
 }
 address 10.127.68.254/32 {
 destination 172.16.74.7;
```

```

 }
 }
}

```

Finally, configure traffic sampling:

```

[edit forwarding-options]
sampling {
 input {
 family inet {
 rate 100;
 run-length 2;
 }
 }
 family inet {
 output {
 file {
 filename sonet-samples.txt;
 files 40;
 size 5m;
 }
 }
 }
}

```

### Example: Sampling All Traffic from a Single IP Address

The following configuration gathers statistical information about every packet entering the router on a specific Gigabit Ethernet port originating from a single source IP address of 172.16.92.31, and collects it in a file named **samples-172-16-92-31.txt**.

Create the filter:

```

[edit firewall family inet]
filter one-ip {
 term get-ip {
 from {
 source-address 172.16.92.31;
 }
 then {
 sample;
 accept;
 }
 }
}

```

Apply the filter to the Gigabit Ethernet interface:

```

[edit interfaces]
ge-4/1/1 {
 unit 0 {
 family inet {
 filter {
 input one-ip;
 }
 address 10.45.92.254;
 }
 }
}

```

```
 }
 }
}
```

Finally, gather statistics on all the candidate samples; in this case, gather all statistics:

```
[edit forwarding-options]
sampling {
 input {
 family inet {
 rate 1;
 }
 }
 family inet {
 output {
 file {
 filename samples-172-16-92-31.txt;
 files 100;
 size 100k;
 }
 }
 }
}
```

---

### Example: Sampling All FTP Traffic

The following configuration gathers statistical information about a moderate percentage of packets using the FTP data transfer protocol in the output path of a specific T3 interface, and collects the information in a file named **t3-ftp-traffic.txt**.

Create a filter:

```
[edit firewall family inet]
filter ftp-stats {
 term ftp-usage {
 from {
 destination-port [ftp ftp-data];
 }
 then {
 sample;
 accept;
 }
 }
}
```

Apply the filter to the T3 interface:

```
[edit interfaces]
t3-7/0/2 {
 unit 0 {
 family inet {
 filter {
 input ftp-stats;
 }
 address 10.35.78.254/32 {
 destination 10.35.78.4;
 }
 }
 }
}
```

```

 }
 }
}

```

Finally, gather statistics on 10 percent of the candidate samples:

```

[edit forwarding-options]
sampling {
 input {
 family inet {
 rate 10;
 }
 }
 family inet {
 output {
 file {
 filename t3-ftp-traffic.txt;
 files 50;
 size 1m;
 }
 }
 }
}

```

- Related Documentation**
- *Traffic Sampling, Forwarding, and Monitoring Overview*
  - [Sampling Instance Configuration on page 881](#)

## Sampling Instance Configuration

You can configure active sampling by defining a sampling instance that specifies a name for the sampling parameters and bind the instance name to an FPC, MPC, or DPC. This configuration enables you to define multiple named sampling parameter sets associated with multiple destinations and protocol families per sampling destination. With the cflowd version 5 and version 8 and flow aggregation version 9, you can use templates to organize the data gathered from sampling.

To implement this feature, you include the **instance** statement at the **[edit forwarding-options sampling]** hierarchy level.

The following considerations apply to the sampling instance configuration:

- This configuration is supported on the IP version 4 (**inet**), IP version 6 (**ipv6**), and MPLS protocol families.
- You can configure the router to perform sampling in either of two locations:
  - On the Routing Engine, using the sampled process. To select this method, use a filter (input or output) with a matching term that contains the **then** sample statement.
  - On the Monitoring Services, Adaptive Services, or Multiservices PIC. Specify the interface name at the **[forwarding-options sampling instance *instance-name* family inet output interface]** hierarchy level. You can configure the same or different services PICs in a set of sampling instances.

- You can configure the **rate** and **run-length** options at the **[edit forwarding-options sampling input]** hierarchy level to apply common values for all families on a global basis. Alternatively, you can configure these options at the **[edit forwarding-options sampling instance *instance-name* input]** hierarchy level to apply specific values for each instance or at the **[edit forwarding-options sampling instance *instance-name* family *family* input]** hierarchy level to apply specific values for each protocol family you configure.
- For MX Series devices with Modular Port Concentrators (MPCs), port-mirrored or sampled packets can be truncated (or clipped) to any length in the range of 1 through 255 bytes. Only the values 1 to 255 are valid for packet truncation on these devices. For other devices, the range is from 0 through 9216. A maximum-packet-length value of zero (0) represents that truncation is disabled, and the entire packet is mirrored or sampled.



**NOTE:** The **run-length** and **maximum-packet-length** configuration statements are not supported on MX80 routers.

---

To associate the defined instance with a particular FPC, MPC, or DPC, you include the **sampling-instance** statement at the **[edit chassis fpc *number*]** hierarchy level, as in the following example:

```
chassis {
 fpc 2 {
 sampling-instance samp1;
 }
}
```

To associate a sampling instance with an FPC in the MX Series Virtual Chassis master or backup router, use the **sampling-instance *instance-name*** statement at the **[edit chassis member *member-number* fpc slot *slot-number*]** hierarchy level, where *member-number* is 0 (for the master router) or 1 (for the backup router), and *slot-number* is a number in the range 0 through 11.

#### Related Documentation

- *Traffic Sampling, Forwarding, and Monitoring Overview*
- *Flow Monitoring Feature Guide for Routing Devices*
- *More Information About Flow Monitoring*
- *Configuring Active Flow Monitoring*
- *Directing Traffic Sampling Output to a Server Running the cflowd Application*
- [Configuring Traffic Sampling on page 871](#)
- *Example: Sampling Instance Configuration*
- *[edit forwarding-options sampling] Hierarchy Level*
- *Inline Flow Monitoring for Virtual Chassis Overview*



## Configuring Discard Accounting

---

Discard accounting is similar to traffic sampling, but varies from it in two ways:

- In discard accounting, the packet is intercepted by the monitoring PIC and is not forwarded to its destination.
- Traffic sampling allows you to limit the number of packets sampled by configuring the **max-packets-per-second**, **rate**, and **run-length** statements. Discard accounting does not provide these options, and a high packet count can potentially overwhelm the monitoring PIC.

A discard instance is a named entity that specifies collector information under the **accounting name** statement. Discard instances are referenced in firewall filter **term** statements by including the **then discard accounting name** statement.

Most of the other statements are also found at the **[edit forwarding-options sampling]** hierarchy level. For information on cflowd, see [“Enabling Flow Aggregation” on page 898](#). The **flow-active-timeout** and **flow-inactive-timeout** statements are described in [“Configuring Flow Monitoring” on page 818](#).

To direct sampled traffic to a flow-monitoring interface, include the **interface** statement. The **engine-id** and **engine-type** statements specify the accounting interface used on the traffic, and the **source-address** statement specifies the traffic source.

You cannot use rate-limiting with discard accounting; however, you can specify the duration of the interval for exporting aggregated accounting information by including the **aggregate-export-interval** statement in the configuration. This enables you to put a boundary on the amount of traffic exported to a flow-monitoring interface.

- Related Documentation**
- [Enabling Flow Aggregation on page 898](#)
  - [Configuring Flow Monitoring on page 818](#)



# Sampling Data Using Inline Sampling

- [Understanding Inline Active Flow Monitoring on page 1065](#)
- [Configuring Inline Active flow Monitoring on page 1070](#)
- [Configuring Inline Active Flow Monitoring on MX80 Routers on page 1074](#)

## Understanding Inline Active Flow Monitoring

---

This topic provides an overview of the inline active flow monitoring feature and IPFIX and Version 9 flow collection templates used for inline active flow monitoring.

This topic contains the following sections:

- [Inline Active Flow Monitoring on page 1065](#)
- [Inline Active Flow Monitoring Limitations and Restrictions on page 1066](#)
- [IPFIX and Version 9 Templates on page 1067](#)

## Inline Active Flow Monitoring

The inline active flow monitoring is implemented on the Packet Forwarding Engine. All the functions like flow creation, flow update, and flow records export are done by the Packet Forwarding Engine. The flow records are sent out in industry standard IPFIX format.

Inline active flow monitoring provides for higher scalability and performance as the scaling and performance are not dependent on the capacity of the services interface. It is also cost effective in more than one way as there is no need to invest in additional hardware or to dedicate a PIC slot for the services PIC. You can make full use of the available slots for handling traffic on the device.

Junos OS Release 13.2 extends inline active flow monitoring support to VPLS flows. Now, you can configure inline active flow monitoring for IPv4, IPv6, and VPLS traffic.

The inline active flow monitoring configuration can be broadly classified into four categories:

1. Configurations at the **[edit services flow-monitoring]** hierarchy level—At this level, you configure the template properties for inline flow monitoring.
2. Configurations at the **[edit forwarding-options]** hierarchy level—At this level, you configure a sampling instance and associate the template (configured at the **[edit**

**services flow-monitoring**] hierarchy level) with the sampling instance. At this level, you also configure the flow-server IP address and port number as well as the flow export rate.

3. Configurations at the **[edit chassis]** hierarchy level—At this level, you associate the sampling instance with the FPC on which the media interface is present. If you are configuring sampling of IPv6 flows, you must also specify the flow hash table size.
4. Configurations at the **[edit firewall]** hierarchy level—At this level you configure a firewall filter for the family of traffic to be sampled. You must attach this filter to the interface on which you want to sample the traffic.

Inline active flow monitoring supports version 9 and IPFIX flow collection templates. Support for version 9 template was introduced in Junos OS Release 13.2, and is limited to IPv4 flows. IPFIX template is supported for IPv4, IPv6, and VPLS flows. IPFIX template uses UDP as the transport protocol, whereas version 9 is transport protocol-independent.

Before you configure inline active flow monitoring, you should ensure that you have adequately-sized hash tables for IPv4 and IPv6 flow sampling. These tables can use one to fifteen 256k areas, and each table is assigned a default value of one such area. When anticipated traffic volume requires larger tables, allocate larger tables.



**NOTE:** Starting with Junos OS Release 13.3, you can configure flow collectors to be reachable through non-default VPN routing and forwarding (VRF) instances by including the `routing-instance instance-name` statement at the `[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-server hostname]` hierarchy level for inline flow monitoring. You cannot configure a flow collector to be reachable through non-default VRF instances for version 5 and version 8 flows. You must configure the routing instance to be a VRF instance by including the `instance-type vrf` statement at the `[edit routing-instances instance-name]` hierarchy level.

---

## Inline Active Flow Monitoring Limitations and Restrictions

The following limitations and restrictions apply to the inline active flow monitoring feature in Junos OS:

- You can configure inline active flow monitoring only on MX Series routers with Trio-based line cards and T4000 routers with Type 5 FPCs.
- You can apply Version 9 flow template only to IPv4 traffic.
- You can configure only one sampling instance on an Flexible PIC Concentrator (FPC).
- You can configure only one type of sampling—either PIC-based sampling or inline sampling—per family in a sampling instance. However, you can configure PIC-based and inline sampling for different families in a sampling instance.
- You can configure only one collector for inline active flow monitoring.

- The following considerations apply to the inline flow-monitoring instance configuration:
  - Sampling run-length and clip-size are not supported.
  - For inline configurations, each family can support only one collector.
  - The user-defined sampling instance gets precedence over the global instance. When a user-defined sampling instance is attached to the FPC, the global instance is removed from the FPC and the user-defined sampling instance is applied to the FPC.
- On routers with Multiservices PICs or Multiservices DPCs, all fragments of a fragmented IPv4 packet other than the first fragment of the packet are processed accurately by the flow monitoring application running on MS-PIC or MS-DPC. The flow monitoring mechanism handles such fragments accurately by setting the layer 4 related fields in the associated flows to zero.
- Flow records and templates cannot be exported if the flow collector is reachable through any management interface.
- The flow collector should be reachable through the default routing table (inet.0 or inet6.0). If the flow collector is reachable via a non-default VPN routing and forwarding table (VRF), flow records and templates cannot be exported.



**NOTE:** Starting with Junos OS Release 13.3, you can configure the flow collector to be reachable through non-default VRF instances apart from being reachable over the default VRF instance. Flow records and templates can be exported even with non-default VRF instances.

- If the destination of the sampled flow is reachable through multiple paths, the IP\_NEXT\_HOP (Element ID 15) and OUTPUT\_SNMP (Element ID 14) in the IPv4 flow record would be set to the Gateway Address and SNMP Index of the first path seen in the forwarding table.
- If the destination of the sampled flow is reachable through multiple paths, the IP\_NEXT\_HOP (Element ID 15) and OUTPUT\_SNMP (Element ID 14) in the IPv6 flow records would be set to 0.
- The Incoming Interface (IIF) and Outgoing Interface (OIF) should be part of the same VRF. If OIF is in a different VRF, DST\_MASK (Element ID 13), DST\_AS (Element ID 17), IP\_NEXT\_HOP (Element ID 15), and OUTPUT\_SNMP (Element ID 14) would be set to 0 in the flow records.
- Each Lookup Chip (LU) maintains and exports flows independent of other LUs. Traffic received on a media interface is distributed across all LUs in a multi-LU platform. It is likely that a single flow will be processed by multiple LUs. Therefore, each LU creates a unique flow and exports it to the flow collector. This can cause duplicate flows records to be seen on the flow collector. The flow collector should aggregate PKTS\_COUNT and BYTES\_COUNT for duplicate flow records to derive a single flow record.

## IPFIX and Version 9 Templates

The following sections list the fields included in IPFIX and Version 9 templates.

### Fields Included in the IPFIX IPv4 Template

---

- IPv4 Source Address
- IPv4 Destination Address
- IPv4 TOS
- IPv4 Protocol
- L4 Source Port
- L4 Destination Port
- ICMP Type and Code
- Input Interface
- VLAN ID
- IPv4 Source Mask
- IPv4 Destination Mask
- Source AS
- Destination AS
- IPv4 Next Hop Address
- TCP Flags
- Output Interface
- Number of Flow Bytes
- Number of Flow Packets
- Minimum TTL (time to live)
- Maximum TTL (time to live)
- Flow Start Time
- Flow End Time
- Flow End Reason
- 802.1Q VLAN identifier (dot1qVlanId)
- 802.1Q Customer VLAN identifier (dot1qCustomerVlanId)

### Fields Included in the IPFIX IPv6 Template

---

- IPv6 Source Address
- IPv6 Destination Address
- IPv6 TOS
- IPv6 Protocol
- L4 Source Port

- L4 Destination Port
- ICMP Type and Code
- Input Interface
- VLAN ID
- IPv6 Source Mask
- IPv6 Destination Mask
- Source AS
- Destination AS
- IPv6 Next Hop Address
- TCP Flags
- Output Interface
- Number of Flow Bytes
- Number of Flow Packets
- Minimum Hop Limits
- Maximum Hop Limits
- Flow Start Time
- Flow End Time
- Flow End Reason
- 802.1Q VLAN identifier (dot1qVlanId)
- 802.1Q Customer VLAN identifier (dot1qCustomerVlanId)

---

#### Fields Included in the Version 9 IPv4 Template

- IPv4 Source Address
- IPv4 Destination Address
- IPv4 TOS
- IPv4 Protocol
- L4 Source Port
- L4 Destination Port
- ICMP Type and Code
- Input Interface
- VLAN ID
- IPv4 Source Mask
- IPv4 Destination Mask
- Source AS

- Destination AS
- IPv4 Next Hop Address
- BGP IPv4 Next Hop Address
- TCP Flags
- Output Interface
- Number of Flow Bytes
- Number of Flow Packets
- Time when the first packet of the flow was switched.
- Time when the last packet of flow was switched.
- Internet Protocol Version

**Related  
Documentation**

- *Example: Configuring Inline Active Flow Monitoring*
- [Configuring Inline Active Flow Monitoring on MX80 Routers on page 894](#)

---

## Configuring Inline Active flow Monitoring

The inline active flow monitoring is implemented on the Packet Forwarding Engine. All the functions like flow creation, flow update, and flow records export are done by the Packet Forwarding Engine. The flow records are sent out in industry standard IPFIX format.

The inline active flow monitoring configuration can be broadly classified into four categories:

1. Configurations at the **[edit services flow-monitoring]** hierarchy level—At this level, you configure the template properties for inline flow monitoring.
2. Configurations at the **[edit forwarding-options]** hierarchy level—At this level, you configure a sampling instance and associate the template (configured at the **[edit services flow-monitoring]** hierarchy level) with the sampling instance. At this level, you also configure the flow-server IP address and port number as well as the flow export rate.
3. Configurations at the **[edit chassis]** hierarchy level—At this level, you associate the sampling instance with the FPC on which the media interface is present. If you are configuring sampling of IPv6 flows, you must also specify the flow hash table size.
4. Configurations at the **[edit firewall]** hierarchy level—At this level you configure a firewall filter for the family of traffic to be sampled. You must attach this filter to the interface on which you want to sample the traffic.

Before you configure inline active flow monitoring, you should ensure that you have adequately-sized hash tables for IPv4 and IPv6 flow sampling. These tables can use one to fifteen 256k areas, and each table is assigned a default value of one such area. When anticipated traffic volume requires larger tables, allocate larger tables.





**NOTE:** For Junos OS releases earlier than Release 12.1, the following points are applicable for supporting backward compatibility when you configure the IPv4 and IPv6 flow table sizes for inline active flow monitoring:

- If you do not configure the `flow-table-size` statement at the `[edit chassis fpc slot-number inline-services]` hierarchy level, fifteen 256K entries are allocated by default for the IPv4 flow table and one 1K entry is allocated by default for the IPv6 flow table on the Packet Forwarding Engine.
- If you configure the `ipv4-flow-table-size size` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level and if you do not configure the `ipv6-flow-table-size size` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level, the number of units of 256K entries that you configure for the IPv4 flow table is allocated. For the IPv6 flow table, a default size of one 1K entry is allocated on the Packet Forwarding Engine.
- If you do not configure the `ipv4-flow-table-size size` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level and if you configure the `ipv6-flow-table-size size` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level, the number of units of 256K entries that you configure for the IPv6 flow table is allocated. For the IPv4 flow table, a default size of one 1K entry is allocated on the Packet Forwarding Engine.
- If you configure the sizes of both the IPv4 and IPv6 flow tables, the flow tables are created on the Packet Forwarding Engine based on the size that you specified.



**NOTE:** The functionality to log the cflowd records in a log file before they are exported to a cflowd server (by including the `local-dump` statement at the `[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-server hostname]` hierarchy level) is not supported when you configure inline flow monitoring (by including the `inline-jflow` statement at the `[edit forwarding-options sampling instance instance-name family inet output]` hierarchy level).

To allocate IPv4 and IPv6 flow hash tables:

1. Go to the `flow-table-size` hierarchy level for inline services on the FPC that processes the monitored flows.

```
[edit]
user@host# edit chassis fpc 0 inline-services flow-table-size
```

2. Specify the required sizes for the sampling hash tables.

```
[edit chassis fpc 0 inline-services flow-table-size]
user@host# set ipv4-flow-table-size 5
user@host# set ipv6-flow-table-size 5
```



**NOTE:** When you set the flow hash table sizes, remember:

- Any change in the configured size of flow hash table sizes initiates an automatic reboot of the FPC.
- The total number of units used for both IPv4 and IPv6 cannot exceed 15.

To configure inline active flow monitoring on all other MX Series routers (except for MX80 routers), EX Series switches, and T4000 routers with Type 5 FPC:

1. Enable inline active flow monitoring and specify the source address for the traffic.

```
[edit forwarding-options sampling instance instance-name family inet output]
user@host# set inline-jflow source address address
```

2. Specify the IP\_FIX output format.

```
[edit forwarding-options sampling instance instance-name family inet output flow-server
address]
user@host# set version-ipfix template ipv4
```

3. Specify the output properties.

```
[edit services flow-monitoring]
user@host# set version-ipfix
```

The output format properties are common to other output formats and are described in [“Configuring Flow Aggregation to Use IPFIX Flow Templates” on page 912](#).

The following is an example of the sampling configuration for an instance that supports inline active flow monitoring on **family inet** and PIC-based sampling on **family inet6**:

```
[edit forwarding-options]
sampling {
 instance {
 sample-ins1 {
 input {
 rate 1;
 }
 family inet {
 output {
 flow-server 2.2.2.2 {
 port 2055;
 version-ipfix {
 template {
 ipv4;
 }
 }
 }
 }
 inline-jflow {
 source-address 10.11.12.13;
 }
 }
 }
 }
}
```

```

family inet6 {
 output {
 flow-server 2.2.2.2 {
 port 2055;
 version-ipfix {
 template {
 ipv6;
 }
 }
 }
 }
 interface sp-0/1/0 {
 source-address 10.11.12.13;
 }
}
}
}

```

The following example shows the output format configuration:

```

services {
 flow-monitoring {
 version-ipfix {
 template ipv4 {
 flow-active-timeout 60;
 flow-inactive-timeout 60;
 ipv4-template;
 template-refresh-rate {
 packets 1000;
 seconds 10;
 }
 }
 option-refresh-rate {
 packets 1000;
 seconds 10;
 }
 }
 }
}
}

```

The following considerations apply to the inline flow-monitoring instance configuration:

- Sampling run-length and clip-size are not supported.
- For inline configurations, each family can support only one collector.



**NOTE:** On routers with Multiservices PICs or Multiservices DPCs, IPv4 and IPv6 fragments are processed accurately. The flow monitoring application creates two flows for every fragmented flow. The first fragment that has the complete Layer 4 information forms the first flow with 5-tuple data and subsequently, all the fragmented packets related to this flow form another flow with the Layer 4 fields set to zero.

- Related Documentation**
- [Configuring Inline Active Flow Monitoring on MX80 Routers on page 894](#)
  - [inline-jflow on page 1687](#)

---

## Configuring Inline Active Flow Monitoring on MX80 Routers

---

To configure inline active flow monitoring on MX80 routers:

1. Associate a sampling instance with the Forwarding Engine Processor.

```
[edit]
user@host# set chassis tfeb slot number sampling-instance sampling-instance
```

The Forwarding Engine Processor slot is always 0 because MX80 routers have only one Packet Forwarding Engine. In this configuration, the sampling instance is **sample-ins1**.

```
[edit]
user@host# set chassis tfeb slot 0 sampling-instance sample-ins1
```



**NOTE:** MX80 routers support only one sampling instance.

2. Under forwarding-options, configure a sampling instance for the flow server and inline jflow instances (these will be configured in the following steps):

```
[edit forwarding-options sampling]
user@host# edit instance inline_sample
```

3. Configure the rate at the **[edit forwarding-options sampling instance instance-name input]** hierarchy level to apply specific values for the sampling instance **sample-ins1**.

```
[edit forwarding-options sampling instance sample-ins1 input]
user@host# set rate number
```

In this configuration, the rate is 1000.

```
[edit forwarding-options sampling instance sample-ins1 input]
user@host# set rate 1000
```

4. Navigate to the output hierarchy and from there, enable a flow server and then specify the output address and port:

```
[edit] forwarding-options sampling instance inline_sample family inet output]
user@host# edit flow-server address
```

```
[edit forwarding-options sampling instance inline_sample family inet output flow-server
<address>]
user@host# set port number
```

5. Return to the output hierarchy and specify the source address for inline jflow:

```
[edit forwarding-options sampling instance sample-ins1 family inet output]
user@host# set inline-jflow source-address address
```

In this configuration, the source address is 10.11.12.13.

```
[edit forwarding-options sampling instance sample-ins1 family inet output]
```

```
user@host# set inline-jflow source-address 10.11.12.13
```

6. Specify the output properties.

```
[edit services flow-monitoring]
```

```
user@host# set version-ipfix
```

The output format properties are common to other output formats and are described in [“Configuring Flow Aggregation to Use IPFIX Flow Templates” on page 912](#).

The following is an example of the sampling configuration for an instance that supports inline active flow monitoring on MX80 routers:

```
[edit forwarding-options]
user@host# show
sampling {
 instance {
 sample-ins1 {
 input {
 rate 1000;
 }
 family inet {
 flow-server 133..13.13.122{
 port 1333;
 inline-jflow {
 source-address 10.11.12.13;
 }
 }
 }
 }
 }
}
```



**NOTE:** You need not configure a Flexible PIC Concentrator (FPC) slot because MX80 routers have only one Packet Forwarding Engine.

The following considerations apply to the inline flow-monitoring instance configuration:

- This configuration does not support MPLS-IPv6.
- Clip-size is not supported.

#### Related Documentation

- [Configuring Flow Aggregation to Use IPFIX Flow Templates on page 912](#)
- [Configuring Inline Active flow Monitoring on page 890](#)
- [inline-jflow on page 1687](#)



# Sampling Data Using Flow Aggregation

- Understanding Flow Aggregation on page 1077
- Enabling Flow Aggregation on page 1078
- Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd on page 1078
- Configuring Flow Aggregation to Use Version 9 Flow Templates on page 1082
- Configuring Flow Aggregation to Use IPFIX Flow Templates on page 1092
- Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows on page 1098
- Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows on page 1101
- Directing Replicated Flows to Multiple Flow Servers on page 1106
- Logging cflowd Flows Before Export on page 1108

## Understanding Flow Aggregation

---

You can collect an aggregate of sampled flows and send the aggregate to a specified host that runs either the cflowd application available from CAIDA (<http://www.caida.org>) or the newer version 9 format defined in RFC 3954, *Cisco Systems NetFlow Services Export Version 9*. Before you can perform flow aggregation, the routing protocol process must export the autonomous system (AS) path and routing information to the sampling process.

By using flow aggregation, you can obtain various types of byte and packet counts of flows through a router. The application collects the sampled flows over a period of 1 minute. At the end of the minute, the number of samples to be exported are divided over the period of another minute and are exported over the course of the same minute.

You configure flow aggregation in different ways, depending on whether you want to export flow records in cflowd version 5 or 8 format, or the separate version 9 format. The latter allows you to sample MPLS, IPv4, IPv6, and peer AS billing traffic. You can also combine configuration statements between the MPLS and IPv4 formats.



**NOTE:** When PIC-based sampling is enabled, collection of flow statistics for sampled packets on flows in virtual private networks (VPNs) is also supported. No additional CLI configuration is required.

- Related Documentation**
- [Enabling Flow Aggregation on page 898](#)
  - [Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd on page 898](#)
  - [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902](#)
  - [Directing Replicated Flows to Multiple Flow Servers on page 926](#)
  - [Logging cflowd Flows Before Export on page 928](#)

---

## Enabling Flow Aggregation

Before you can perform flow aggregation, the routing protocol process must export the autonomous system (AS) path and routing information to the sampling process. To enable the export of AS path and the routing information to the sampling process, one or more of the following needs to be configured:

- At the `[edit forwarding-options]` hierarchy level (for routing instances, at the `[edit routing-instance routing-instance-name forwarding-options]` hierarchy level), configure `sampling family` or `sampling output` or `sampling instance` or `monitoring` or `accounting`.
- At the `[edit routing-options]` hierarchy level (for routing instances, at the `[edit routing-instance routing-instance-name routing-options]` hierarchy level), configure `route record`.
- At the `[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]` hierarchy level, configure `forwarding-db-size`.

- Related Documentation**
- [Understanding Flow Aggregation on page 897](#)
  - [Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd on page 898](#)
  - [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902](#)
  - [Directing Replicated Flows to Multiple Flow Servers on page 926](#)
  - [Configuring Traffic Sampling on page 871](#)
  - [Example: Configuring Active Flow Monitoring Version 9 for IPv6](#)
  - [Logging cflowd Flows Before Export on page 928](#)

---

## Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd

To enable the collection of cflowd version 5 or version 8 flow formats, include the `flow-server` statement:

```
flow-server hostname {
 aggregation {
 autonomous-system;
 destination-prefix;
 protocol-port;
 source-destination-prefix {
 caida-compliant;
 }
 }
}
```



```

 }
 source-prefix;
 }
 autonomous-system-type (origin | peer);
 (local-dump | no-local-dump);
 port port-number;
 version format;
}

```

You can include this statement at the following hierarchy levels:

- [edit forwarding-options sampling family (inet | inet6 | mpls) output]
- [edit forwarding-options sampling instance *instance-name* output]
- [edit forwarding-options accounting *name* output cflowd *hostname*]

You must configure the **family inet** statement on logical interface **unit 0** on the monitoring interface, as in the following example:

```

[edit interfaces]
sp-3/0/0 {
 unit 0 {
 family inet {
 ...
 }
 }
}

```



**NOTE:** Boot images for monitoring services interfaces are specified at the [edit chassis images pic] hierarchy level. You must enable the NTP client to make the cflowd feature operable, by including the following configuration:

```

[edit system]
ntp {
 boot-server ntp.juniper.net;
 server 172.17.28.5;
}
processes {
 ntp enable;
}

```

For more information, see the *Junos OS Administration Library for Routing Devices*.

You can also configure cflowd version 5 for flow-monitoring applications by including the **cflowd** statement at the [edit forwarding-options monitoring *name* family inet output] hierarchy level:

```

cflowd hostname {
 port port-number;
}

```

The following restrictions apply to cflowd flow formats:

- You can configure up to one version 5 and one version 8 flow format at the **[edit forwarding-options accounting *name* output]** hierarchy level.
- You can configure up to eight version 5 or one version 8 flow format at the **[edit forwarding-options sampling family (inet | inet6 | mpls) output]** hierarchy level for Routing Engine-based sampling by including the **flow-server** statement. In contrast, PIC-based sampling allows you to specify one cflowd version 5 server and one version 8 server simultaneously. However, the two cflowd servers must have different IP addresses.
- You can configure up to eight version 5 flow formats at the **[edit forwarding-options monitoring *name* output]** hierarchy level. Version 8 flow formats and aggregation are not supported for flow-monitoring applications.
- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created on the monitoring PIC only after the route record resynchronization operation is complete, which is 60 seconds after the PIC comes up. Any packets sent to the PIC would be dropped until the synchronization process is complete.
- The configuration includes a proprietary v5 extension template for supporting 4-byte AS information in flow records. Its template version is set to 500, indicating it to be proprietary. All other fields remain the same; the source AS and destination AS are each 4 bytes long, rather than 2 bytes as in the traditional v5 template. This option is available at the **[edit forwarding-options sampling family inet output flow-server server-name version]** hierarchy level.

In the **cflowd** statement, specify the name or identifier of the host that collects the flow aggregates. You must also include the User Datagram Protocol (UDP) port number on the host and the version, which gives the format of the exported cflowd aggregates. To collect cflowd records in a log file before exporting, include the **local-dump** statement.



**NOTE:** You can specify both host (cflowd) sampling and port mirroring in the same configuration; however, only one action takes effect at any one time. Port mirroring takes precedence. For more information, see [“Configuring Port Mirroring” on page 931](#).

For cflowd version 8 only, you can specify aggregation of specific types of traffic by including the **aggregation** statement. This conserves memory and bandwidth by enabling cflowd to export targeted flows rather than all aggregated traffic. To specify a flow type, include the **aggregation** statement:

```
aggregation {
 autonomous-system;
 destination-prefix;
 protocol-port;
 source-destination-prefix {
 caida-compliant;
 }
}
```

```

 }
 source-prefix;
 }

```

You can include this statement at the following hierarchy levels:

- [edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server *hostname*]
- [edit forwarding-options accounting *name* output cflowd *hostname*]

The **autonomous-system** statement configures aggregation by the AS number; this statement might require setting the separate cflowd **autonomous-system-type** statement to include either **origin** or **peer** AS numbers. The **origin** option specifies to use the origin AS of the packet source address in the Source Autonomous System cflowd field. The **peer** option specifies to use the peer AS through which the packet passed in the Source Autonomous System cflowd field. By default, cflowd exports the origin AS number.

The **destination-prefix** statement configures aggregation by the destination prefix only.

The **protocol-port** statement configures aggregation by the protocol and port number; requires setting the separate **cflowd port** statement.

The **source-destination-prefix** statement configures aggregation by the source and destination prefix. Version 2.1b1 of CAIDA's cflowd application does not record source and destination mask length values in compliance with CAIDA's *cflowd Configuration Guide*, dated August 30, 1999. If you configure the **caida-compliant** statement, the Junos OS complies with Version 2.1b1 of cflowd. If you do not include the **caida-compliant** statement in the configuration, the Junos OS records source and destination mask length values in compliance with the *cflowd Configuration Guide*.

The **source-prefix** statement configures aggregation by the source prefix only.

Collection of sampled packets in a local ASCII file is not affected by the **cflowd** statement.

The following commands enable RE- and PIC-based sampling at the **set forwarding options sampling** hierarchy level:

- set input rate *rate*
- set input run-length *length*
- set family inet output flow-server *flowcollector* port *udp port*
- set family inet output flow-server *flowcollector* no-local-dump
- set family inet output flow-server *flowcollector* version <5/8>

The following commands enable RE- and PIC-based sampling at the **set interfaces** hierarchy level:

- *interface to be sampled* unit *unit* family inet filter *input/output filename*

The following commands enable RE- and PIC-based sampling at the **set firewall family** hierarchy level:

- `set inet filter filtername term 1 then count filternameing`
- `set inet filter filtername term 1 then sample`
- `set inet filter filtername term 1 then accept`

The following command enables PIC-based sampling at the **set forwarding options sampling** hierarchy level:

- `set family inet output interface sp-*/*/* source address source address`

The following example shows a PIC-based flow aggregation configuration using version 5:

```
family inet {
 output {
 flow-inactive-timeout 15;
 flow-active-timeout 60;
 flow-server 153.104.248.37 {
 port 9996;
 version 5;
 }
 interface sp-2/2/0 {
 engine-id 4;
 source-address 153.104.0.254;
 }
 }
}
```

The following example shows an RE-based flow aggregation configuration using version 5:

```
family inet {
 output {
 flow-inactive-timeout 15;
 flow-active-timeout 60;
 flow-server 153.104.248.37 {
 port 9996;
 source-address 153.104.0.254;
 version 5;
 }
 }
}
```

#### Related Documentation

- [Understanding Flow Aggregation on page 897](#)
- [Enabling Flow Aggregation on page 898](#)
- [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902](#)
- [Configuring Flow Aggregation to Use IPFIX Flow Templates on page 912](#)

---

## Configuring Flow Aggregation to Use Version 9 Flow Templates

Use of version 9 allows you to define a flow record template suitable for IPv4 traffic, IPv6 traffic, MPLS traffic, a combination of IPv4 and MPLS traffic, or peer AS billing traffic.

Templates and the fields included in the template are transmitted to the collector periodically, and the collector need not be aware of the router configuration.



**NOTE:** Version 9 requires that you install a services PIC, such as the Adaptive Services PIC or Multiservices PIC in the router. On MX Series routers, the Multiservices DPC fulfills this requirement. For more information on determining which services PIC is suitable for your router, see [“Enabling Service Packages” on page 11](#) or the appropriate hardware documentation.



**NOTE:** If multiple protocol families are configured for a particular flow collector, the export packets will originate from multiple Source IDs, with each Source ID corresponding to a particular protocol. The multiple Source IDs do not indicate that the export packets are originating from multiple Service PICs.

The following sections contain additional information:

- [Configuring the Traffic to Be Sampled on page 1083](#)
- [Configuring the Version 9 Template Properties on page 1084](#)
- [Customizing Template ID, Observation Domain ID, and Source ID for Version 9 flow Templates on page 1085](#)
- [Restrictions on page 1086](#)
- [Fields Included in Each Template Type on page 1086](#)
- [MPLS Sampling Behavior on page 1088](#)
- [Verification on page 1088](#)
- [Examples: Configuring Version 9 Flow Templates on page 1089](#)

## Configuring the Traffic to Be Sampled

To specify sampling of IPv4, IPv6, MPLS, or peer AS billing traffic, include the appropriate configuration of the **family** statement at the **[edit forwarding-options sampling]** hierarchy level:

```
[edit forwarding-options]
sampling {
 family (inet | inet6 | mpls);
}
```

You can include **family inet**, **family inet6**, or **family mpls**.



**NOTE:** If you specify sampling for peer AS billing traffic, the **family** statement supports only IPv4 and IPv6 traffic (inet or inet6). Peer AS billing traffic is enabled only at the global instance hierarchy level and is not available for per Packet Forwarding Engine instances.

After you specify the family of traffic to be sampled, configure the sampling parameters such as the maximum packet length (beyond which the packets are truncated), maximum packets to be sampled per second (beyond which the packets are dropped), the rate (for example, if you specify 10, every 10th packet is sampled), and run length (which specify the number of packets to be sampled after the trigger; that is if the **rate** is set to 10 and **run-length** to 5, five packets starting the 10th packet are sampled).

```
[edit forwarding-options sampling]
input {
 maximum-packet-length bytes
 max-packets-per-second number;
 rate number;
 run-length number;
}
```

## Configuring the Version 9 Template Properties

To define the version 9 templates, include the following statements at the **[edit services flow-monitoring version9]** hierarchy level:

```
[edit services flow-monitoring version9]
template name {
 options-template-id
 template-id
 source-id
 flow-active-timeout seconds;
 flow-inactive-timeout seconds;
 option-refresh-rate packets packets seconds seconds;
 template-refresh-rate packets packets seconds seconds;
 (ipv4-template | ipv6-template | mpls-ipv4-template | mpls-template |
 peer-as-billing-template) {
 label-position [positions];
 }
}
```

The following details apply to the configuration statements:

- You assign each template a unique name by including the **template name** statement.
- You then specify each template for the appropriate type of traffic by including the **ipv4-template**, **ipv6-template**, **mpls-ipv4-template**, **mpls-template**, or **peer-as-billing-template**.
- If the template is used for MPLS traffic, you can also specify up to three label positions for the MPLS header label data by including the **label-position** statement; the default values are **[1 2 3]**.
- Within the template definition, you can optionally include values for the **flow-active-timeout** and **flow-inactive-timeout** statements. These statements have specific default and range values when they are used in template definitions; the default is 60 seconds and the range is from 10 through 600 seconds. Values you specify in template definitions override the global timeout values configured at the **[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server]** hierarchy level.

- You can also include settings for the **option-refresh-rate** and **template-refresh-rate** statements within a template definition. For both of these properties, you can include a timer value (in seconds) or a packet count (in number of packets). For the **seconds** option, the default value is 60 and the range is from 10 through 600. For the **packets** option, the default value is 4800 and the range is from 1 through 480,000.
- To filter IPV6 traffic on a media interface, the following configuration is supported:

```

interfaces interface-name {
 unit 0 {
 family inet6 {
 sampling {
 input;
 output;
 }
 }
 }
}

```

## Customizing Template ID, Observation Domain ID, and Source ID for Version 9 flow Templates

Use of version 9 and IPFIX allows you to define a flow record template suitable for IPv4 traffic, IPv6 traffic, MPLS traffic, a combination of IPv4 and MPLS traffic, or peer AS billing traffic. Templates and the fields included in the template are transmitted to the collector periodically, and the collector need not be aware of the router configuration. Starting with Junos OS Release 14.1, you can specify the unique identifier for the version 9 and IPFIX templates. The identifier of a template is locally unique within a combination of a transport session and an observation domain. Template IDs 0 through 255 are reserved for template sets, options template sets, and other sets for future use. Template IDs of data sets are numbered from 256 through 65535. Typically, this information element or field in the template is used to define the characteristics or properties of other information elements in a template. After a restart of the export process of templates is performed, template IDs can be reassigned. In Junos OS releases earlier than Release 14.1, template IDs and options template IDs were predefined for each address family and could not be modified.

This functionality to configure template ID, options template ID, observation domain ID, and source ID is supported on all routers with MPCs (Trio chip-based FPCs).

The following values were assigned by default for the template IDs of IPFIX templates for the different protocols or address families, until Junos OS Release 13.3:

- IPv4 flow template ID—256
- IPv6 flow template ID—257
- VPLS flow template ID—258
- Options template ID for all address families—512

The corresponding data sets and option data sets contain the value of the template IDs and options template IDs respectively in the set ID field. This method enables the collector to match a data record with a template record.

For more information about specifying the source ID, observation domain ID, template ID, and options template ID for version 9 and IPFIX flows, see [“Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows” on page 918](#) and [“Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows” on page 921](#).

## Restrictions

The following restrictions apply to version 9 templates:

- You cannot apply the two different types of flow aggregation configuration (cflowd version 5/8 and flow aggregation version 9) at the same time.
- Flow export based on an **mpls-ipv4** template assumes that the IPv4 header follows the MPLS header. In the case of Layer 2 VPNs, the packet on the provider router (P router) would look like this:

MPLS | Layer 2 Header | IPv4

In this case, **mpls-ipv4** flows are not created on the PIC, because the IPv4 header does not directly follow the MPLS header. Packets are dropped on the PIC and are accounted as parser errors.

- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created on the monitoring PIC only after the route record resynchronization operation is complete, which is 60 seconds after the PIC comes up. Any packets sent to the PIC would be dropped until the synchronization process is complete.



**NOTE:** "Because the forwarding of a packet that arrives with MPLS labels is performed based on the MPLS label and not based on the IP address contained in the packet, the packet is sampled at the output interface with the MPLS label that was popped not being available at the time of sampling. In such a case, depending on the incoming interface (IIF), the VRF index is identified and the route for the sampled packet is determined in the VRF table. Because a specific route is not available in the VRF that is different from the VRF on which the packet is received, the Output Interface Index, Source Mask, and Destination Mask fields are incorrectly populated. This behavior occurs when an IPv4 template is applied as a firewall filter on an egress interface with sample as the action."

---

## Fields Included in Each Template Type

The following fields are common to all template types:

- Input interface
- Output interface
- Number of bytes



- Number of packets
- Flow start time
- Flow end time

The IPv4 template includes the following specific fields:

- IPv4 Source Address
- IPv4 Destination Address
- L4 Source Port
- L4 Destination Port
- IPv4 TOS
- IPv4 Protocol
- ICMP type and code
- TCP Flags
- IPv4 Next Hop Address
- Source autonomous system (AS) number
- Destination AS number

The IPv6 template includes the following specific fields:

- IPv6 Source Address and Mask
- IPv6 Destination Address and Mask
- L4 Source Port
- L4 Destination Port
- IPv6 TOS
- IPv6 Protocol
- TCP Flags
- IP Protocol Version
- IPv6 Next Hop Address
- Egress Interface Information
- Source Autonomous System (AS) number
- Destination AS number

The MPLS template includes the following specific fields:

- MPLS Label #1
- MPLS Label #2
- MPLS Label #3

- MPLS EXP Information
- FEC IP Address

The MPLS-IPv4 template includes all the fields found in the IPv4 and MPLS templates.

The peer AS billing template includes the following specific fields:

- IPV4 Class of Service (TOS)
- Ingress Interface
- BGP IPV4 Next Hop Address
- BGP Peer Destination AS Number

## MPLS Sampling Behavior

This section describes the behavior when MPLS sampling is used on egress interfaces in various scenarios (label pop or swap) on provider routers (P routers). For more information on configuration and background specific to MPLS applications, see the *MPLS Applications Feature Guide for Routing Devices*.

1. You configure MPLS sampling on an egress interface on the P router and configure an MPLS flow aggregation template. The route action is label *pop* because penultimate hop popping (PHP) is enabled.

Previously, IPv4 packets (only) would have been sent to the PIC for sampling even though you configured MPLS sampling. No flows should be created, with the result that the parser fails.

With the current capability of applying MPLS templates, MPLS flows are created.

2. As in the first case, you configure MPLS sampling on an egress interface on the P router and configure an MPLS flow aggregation template. The route action is label *swap* and the swapped label is 0 (explicit null).

The resulting behavior is that MPLS packets are sent to the PIC. The flow being sampled corresponds to the label before the swap.

3. You configure a Layer 3 VPN network, in which a customer edge router (CE-1) sends traffic to a provider edge router (PE-A), through the P router, to a similar provider edge router (PE-B) and customer edge router (CE-2) on the remote end.

The resulting behavior is that you cannot sample MPLS packets on the PE-A to P router link.

## Verification

To verify the configuration properties, you can use the **show services accounting aggregation template template-name name** operational mode command.

All other **show services accounting** commands also support version 9 templates, except for **show services accounting flow-detail** and **show services accounting aggregation aggregation-type**. For more information about operational mode commands, see the [CLI Explorer](#).

## Examples: Configuring Version 9 Flow Templates

The following is a sample version 9 template configuration:

```
services {
 flow-monitoring {
 version9 {
 template ip-template {
 flow-active-timeout 20;
 flow-inactive-timeout 120;
 ipv4-template;
 }
 template mpls-template-1 {
 mpls-template {
 label-position [1 3 4];
 }
 }
 template mpls-ipv4-template-1 {
 mpls-ipv4-template {
 label-position [1 5 7];
 }
 }
 template peer-as-billing-template-1 {
 peer-as-billing-template;
 }
 }
 }
}
```

The following is a sample firewall filter configuration for MPLS traffic:

```
firewall {
 family mpls {
 filter mpls_sample {
 term default {
 then {
 accept;
 sample;
 }
 }
 }
 }
}
```

The following sample configuration applies the MPLS sampling filter on a networking interface and configures the AS PIC to accept both IPv4 and MPLS traffic:

```
interfaces {
 at-0/1/1 {
 unit 0 {
 family mpls {
 filter {
 input mpls_sample;
 }
 }
 }
 }
}
```

```
 }
 }
 sp-7/0/0 {
 unit 0 {
 family inet;
 family mpls;
 }
 }
}
```

The following example applies the MPLS version 9 template to the sampling output and sends it to the AS PIC:

```
forwarding-options {
 sampling {
 input {
 family mpls {
 rate 1;
 }
 }
 family mpls {
 output {
 flow-active-timeout 60;
 flow-inactive-timeout 30;
 flow-server 1.2.3.4 {
 port 2055;
 version9 {
 template mpls-ipv4-template-1;
 }
 }
 }
 }
 interface sp-7/0/0 {
 source-address 1.1.1.1;
 }
 }
}
```

The following is a sample firewall filter configuration for the peer AS billing traffic:

```
firewall {
 family inet {
 filter peer-as-filter {
 term 0 {
 from {
 destination-class dcu-1;
 interface ge-2/1/0;
 forwarding-class class-1;
 }
 then count count_team_0;
 }
 }
 term 1 {
 from {
 destination-class dcu-2;
 interface ge-2/1/0;
 }
 }
 }
}
```

```

 forwarding-class class-1;
 }
 then count count_team_1;
}
term 2 {
 from {
 destination-class dcu-3;
 interface ge-2/1/0;
 forwarding-class class-1;
 }
 then count count_team_2;
}
}
}
}

```

The following sample configuration applies the peer AS firewall filter as a filter attribute under the forwarding-options hierarchy for CoS-level data traffic usage information collection:

```

forwarding-options {
 family inet {
 filter output peer-as-filter;
 }
}

```

The following sample configuration applies the peer AS DCU policy options to collect usage statistics for the traffic stream for as-path ingressing at a specific input interface with the firewall configuration hierarchy applied as Forwarding Table Filters (FTFs). The configuration functionality with COS capability can be achieved through FTFs for destination-class usage with forwarding-class for specific input interfaces:

```

policy-options {
 policy-statement P1 {
 from {
 protocol bgp;
 neighbor 10.2.25.5; #BGP router configuration;
 as-path AS-1; #AS path configuration;
 }
 then destination-class dcu-1; #Destination class configuration;
 }
 policy-statement P2 {
 from {
 neighbor 1.2.25.5;
 as-path AS-2;
 }
 then destination-class dcu2;
 }
 policy-statement P3 {
 from {
 protocol bgp;
 neighbor 192.2.1.1;
 as-path AS-3;
 }
 then destination-class dcu3;
 }
}

```

```
as-path AS-1 3131:1111:1123;
as-path AS-2 100000;
as-path AS-3 192:29283:2;
}
```

The following example applies the peer-as-billing version 9 template to enable sampling of traffic for billing purposes:

```
forwarding-options {
 sampling {
 }
 input {
 rate 1;
 }
 family inet {
 output {
 flow-server 10.209.15.58 {
 port 300;
 version9 {
 template {
 peer-as;
 }
 }
 }
 }
 interface sp-5/2/0 {
 source-address 2.3.4.5;
 }
 }
}
family inet {
 filter {
 output peer-as-filter;
 }
}
```

**Related  
Documentation**

- [Understanding Flow Aggregation on page 897](#)
- [Enabling Flow Aggregation on page 898](#)
- [Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd on page 898](#)
- [Configuring Flow Aggregation to Use IPFIX Flow Templates on page 912](#)
- [Configuring Traffic Sampling on page 871](#)
- *Example: Configuring Active Flow Monitoring Version 9 for IPv6*

---

## Configuring Flow Aggregation to Use IPFIX Flow Templates

Use of IPFIX allows you to define a flow record template suitable for IPv4 traffic or IPv6 traffic. Templates are transmitted to the collector periodically, and the collector need not be aware of the router configuration. You can define template refresh rate, flow active timeout and inactive timeout.

If flow records are being sent for multiple protocol families (for example, for IPv4 and IPv6), each protocol family flow will have a unique Observation Domain ID.

The following sections contain additional information:

- [Configuring the IPFIX Template Properties on page 1093](#)
- [Restrictions on page 1094](#)
- [Customizing Template ID, Observation Domain ID, and Source ID for IPFIX flow Templates on page 1094](#)
- [Fields Included in the IPv4 Template on page 1095](#)
- [Fields Included in the IPv6 Template on page 1096](#)
- [Verification on page 1096](#)
- [Example: Configuring an IPFIX Flow Templates and Flow Sampling on page 1097](#)

## Configuring the IPFIX Template Properties

To define the IPFIX templates, include the following statements at the **[edit services flow-monitoring version-ipfix]** hierarchy level:

```
[edit services flow-monitoring IPFIX]
template name {
 options-template-id
 template-id
 observation-domain-id
 flow-active-timeout seconds;
 flow-inactive-timeout seconds;
 option-refresh-rate packets packets seconds seconds;
 template-refresh-rate packets packets seconds seconds;
 (ipv4-template | ipv6-template);
}
```

The following details apply to the configuration statements:

- You assign each template a unique name by including the **template *name*** statement.
- You then specify each template for the appropriate type of traffic by including the **ipv4-template** or **ipv6-template**.
- Within the template definition, you can optionally include values for the **flow-active-timeout** and **flow-inactive-timeout** statements. These statements have specific default and range values when they are used in template definitions; the default is 60 seconds and the range is from 10 through 600 seconds.
- You can also include settings for the **option-refresh-rate** and **template-refresh-rate** statements within a template definition. For both of these properties, you can include a timer value (in seconds) or a packet count (in number of packets). For the **seconds** option, the default value is 600 and the range is from 10 through 600. For the **packets** option, the default value is 4800 and the range is from 1 through 480,000.
- To filter IPV6 traffic on a media interface, the following configuration is supported:

```
interfaces interface-name {
 unit 0 {
```

```
family inet6 {
 sampling {
 input;
 output;
 }
}
```

## Restrictions

The following restrictions apply to IPFIX templates:

- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created only after the route record resynchronization operation is complete, which takes 120 seconds.
- VLAN ID field is not valid for egress traffic, and returns a value of 0 for egress traffic.
- The VLAN ID field is updated when a new flow record is created and so, any change in VLAN ID after the record has been created might not be updated in the record.

## Customizing Template ID, Observation Domain ID, and Source ID for IPFIX flow Templates

Use of version 9 and IPFIX allows you to define a flow record template suitable for IPv4 traffic, IPv6 traffic, MPLS traffic, a combination of IPv4 and MPLS traffic, or peer AS billing traffic. Templates and the fields included in the template are transmitted to the collector periodically, and the collector need not be aware of the router configuration. Starting with Junos OS Release 14.1, you can specify the unique identifier for the version 9 and IPFIX templates. The identifier of a template is locally unique within a combination of a transport session and an observation domain. Template IDs 0 through 255 are reserved for template sets, options template sets, and other sets for future use. Template IDs of data sets are numbered from 256 through 65535. Typically, this information element or field in the template is used to define the characteristics or properties of other information elements in a template. After a restart of the export process of templates is performed, template IDs can be reassigned. In Junos OS releases earlier than Release 14.1, template IDs and options template IDs were predefined for each address family and could not be modified.

This functionality to configure template ID, options template ID, observation domain ID, and source ID is supported on all routers with MPCs (Trio chip-based FPCs).

The following values were assigned by default for the template IDs of version 9 templates for the different protocols or address families, until Junos OS Release 13.3:

- IPv4 flow template ID—272
- IPv6 flow template ID—273



- VPLS flow template ID—274
- Options template ID for all address families—520

The corresponding data sets and option data sets contain the value of the template IDs and options template IDs respectively in the set ID field. This method enables the collector to match a data record with a template record.

For more information about specifying the source ID, observation domain ID, template ID, and options template ID for version 9 and IPFIX flows, see [“Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows” on page 918](#) and [“Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows” on page 921](#).

### Fields Included in the IPv4 Template

- IPv4 Source Address
- IPv4 Destination Address
- IPv4 TOS
- IPv4 Protocol
- L4 Source Port
- L4 Destination Port
- ICMP Type and Code
- Input Interface
- VLAN ID
- IPv4 Source Mask
- IPv4 Destination Mask
- Source AS
- Destination AS
- IPv4 Next Hop Address
- TCP Flags
- Output Interface
- Number of Flow Bytes
- Number of Flow Packets
- Minimum TTL (time to live)
- Maximum TTL (time to live)
- Flow Start Time
- Flow End Time
- Flow End Reason

- 802.1Q VLAN identifier (dot1qVlanId)
- 802.1Q Customer VLAN identifier (dot1qCustomerVlanId)

### Fields Included in the IPv6 Template

- IPv6 Source Address
- IPv6 Destination Address
- IPv6 TOS
- IPv6 Protocol
- L4 Source Port
- L4 Destination Port
- ICMP Type and Code
- Input Interface
- VLAN ID
- IPv6 Source Mask
- IPv6 Destination Mask
- Source AS
- Destination AS
- IPv6 Next Hop Address
- TCP Flags
- Output Interface
- Number of Flow Bytes
- Number of Flow Packets
- Minimum Hop Limits
- Maximum Hop Limits
- Flow Start Time
- Flow End Time
- Flow End Reason
- 802.1Q VLAN identifier (dot1qVlanId)
- 802.1Q Customer VLAN identifier (dot1qCustomerVlanId)
- Fragment Identification
- IPv6 Extension Headers

### Verification

The following show commands are supported for IPFIX:

- `show services accounting flow inline-jflow fpc-slot fpc-slot`
- `show services accounting errors inline-jflow fpc-slot fpc-slot`
- `show services accounting status inline-jflow fpc-slot fpc-slot`

### Example: Configuring an IPFIX Flow Templates and Flow Sampling

The following is a sample IPFIX template configuration:

```
services {
 flow-monitoring {
 version-ipfix {
 template ipv4 {
 flow-active-timeout 60;
 flow-inactive-timeout 70;
 template-refresh-rate seconds 30;
 option-refresh-rate seconds 30;
 ipv4-template;
 }
 }
 }
}

chassis;
fpc 0 {
 sampling-instance s1;
}
```

The following example applies the IPFIX template to enable sampling of traffic for billing:

```
forwarding-options {
 sampling {
 instance {
 s1 {
 input {
 rate 10;
 }
 family inet {
 output {
 flow-server 11.11.4.2 {
 port 2055;
 version-ipfix {
 template {
 ipv4;
 }
 }
 }
 }
 inline-jflow {
 source-address 11.11.2.1;
 }
 }
 }
 }
 }
}
```

- Related Documentation**
- [Understanding Flow Aggregation on page 897](#)
  - [Inclusion of Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates](#)
  - [Enabling Flow Aggregation on page 898](#)
  - [Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd on page 898](#)
  - [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902](#)

---

## Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows

For IPFIX flows, an identifier of an Observation Domain is locally unique to an exporting process of the templates. The export process uses the Observation Domain ID to uniquely identify to the collection process in which the flows were metered. We recommend that you configure this ID to be unique for each IPFIX flow. A value of 0 indicates that no specific Observation Domain is identified by this information element. Typically, this attribute is used to limit the scope of other information elements. If the observation domain is not unique, the collector cannot uniquely identify an IPFIX device.

If you configure the same Observation Domain ID for different template types, such as for IPv4 and IPv6, it does not impact flow monitoring because the actual or the base observation domain ID is transmitted in the flow. The actual observation domain ID is derived from the value you configure and also in conjunction with other parameters such as the slot number, lookup chip (LU) instance, Packet Forwarding Engine instance. Such a method of computation of the observation domain ID ensures that this ID is not the same for two IPFIX devices.

Until Junos OS Release 13.3, the observation domain ID is predefined and is set to a fixed value, which is derived from the combination of FPC slot, sampling protocol, PFE Instance and LU Instance fields. This derivation creates a unique observation domain per LU per family. Starting with Junos OS Release 14.1, you can configure the observation domain ID, which causes the first 8 bits of the field to be configured.

The following modifications have been made:

- FPC slots are expanded to 8 bits to enable more slots to be configured in an MX Series Virtual Chassis configuration.
- 8 bits of the configured observation domain ID are used.
- You can configure a value for the observation domain ID in the range of 0 through 255.
- The Protocol field is increased to 3 bits to provide support for additional protocols in inline flow monitoring.
- You can associate the observation domain ID with templates by using the **observation-domain-id *domain-id*** statement at the **[edit services flow- monitoring version-ipfix template *template-name*]** hierarchy level.

For version 9 flows, a 32-bit value that identifies the Exporter Observation Domain is called the source ID. NetFlow collectors use the combination of the source IP address

and the source ID field to separate different export streams originating from the same exporter.

To specify the observation domain ID for IPFIX flows, include the **observation-domain-id domain-id** statement at the **[edit services flow-monitoring version-ipfix template template-name]** hierarchy level.

```
[edit services flow-monitoring version-ipfix]
template template-name {
 observation-domain-id domain-id;
}
```

To specify the source ID for version 9 flows, include the **source-id source-id** statement at the **[edit services flow-monitoring version9 template template-name]** hierarchy level.

```
[edit services flow-monitoring version9]
template template-name {
 source-id source-id;
}
```

Table 32 on page 919 describes observation domain ID values for different combinations of the configured domain ID, protocol family, FPC slot, and the Packet Forwarding Engine and lookup chip instances.

**Table 42: Example of Observation Domain ID**

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id Conf val rsvd 1proto slot LUInst PFEInst  xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
None	IPv4 (0)	1	1	0	0000 0000 0000 1000 0000 0001 0000 0001 0x00080101
None	IPv6 (1)	1	1	0	0000 0000 0000 1001 0000 0001 0000 0001 0x00090101
None	VPLS (2)	1	1	0	0000 0000 0000 1010 0000 0001 0000 0001 0x000A0101

Table 42: Example of Observation Domain ID *(continued)*

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id  Conf val rsvd lproto slot LUInst PFEInst  xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
None	MPLS (3)	1	1	0	0000 0000 0000 1011 0000 0001 0000 0001 0x000B0101
4	IPv4 (0)	1	1	0	0000 0100 0000 1000 0000 0001 0000 0001 0x04080101
190	IPv4 (0)	1	1	0	1101 1110 0000 1000 0000 0001 0000 0001 0xBE080101
4	IPv4 (0)	2	1	1	0000 0100 0000 1000 0000 0010 0001 0001 0x04080211
4	IPv6 (1)	1	1	0	0000 0100 0000 1001 0000 0001 0001 0000 0x04090110
190	IPv6 (1)	1	1	0	1101 1110 0000 1001 0000 0001 0001 0000 0xBE090110
4	VPLS (2)	2	2	0	0000 0100 0000 1010 0000 0010 0010 0000 0x040A0220

Table 42: Example of Observation Domain ID (*continued*)

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id  Conf val rsvd lproto slot LUInst PFEInst  xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
10	IPv4 (0)	28	2	1	0000 1010 0000 1000 0001 1100 0010 0001 0x0A081C21

**Related Documentation** • [Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows on page 921](#)

## Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows

Starting with Junos OS Release 14.1, you can define the template ID for version 9 and IPFIX templates for inline flow monitoring. To specify the template ID for version 9 flows, include the **template-id** *id* statement at the **[edit services flow-monitoring version9 template *template-name*]** hierarchy level.

```
[edit services flow-monitoring version9]
template template-name {
 template-id id;
}
```

To specify the template ID for version IPFIX flows, include the **template-id** statement at the **[edit services flow-monitoring version-ipfix template *template-name*]** hierarchy level.

```
[edit services flow-monitoring version-ipfix]
template template-name {
 template-id id;
}
```

To specify the options template ID for version 9 flows, include the **options-template-id** statement at the **[edit services flow-monitoring version9 template *template-name*]** hierarchy level.

```
[edit services flow-monitoring version9]
template template-name {
 options-template-id id;
}
```

To specify the options template ID for version IPFIX flows, include the **options-template-id** statement at the **[edit services flow-monitoring version-ipfix template *template-name*]** hierarchy level. The template ID and options template ID can be a value in the range of 1024 through 65535.

```
[edit services flow-monitoring version-ipfix]
template template-name {
 options-template-id id;
}
```

The template ID and options template ID can be a value in the range of 1024 through 65535. If you do not configure values for the template ID and options template ID, default values are assumed for these IDs, which are different for the various address families. If you configure the same template ID or options template ID value for different address families, such a setting is not processed properly and might cause unexpected behavior. For example, if you configure the same template ID value for both IPv4 and IPv6, the collector validates the export data based on the template ID value that it last receives. In this case, if IPv6 is configured after IPv4, the value is effective for IPv6 and the default value is used for IPv4.

The following are the default values of template IDs for IPFIX flows for the different protocols or address families, until Junos OS Release 13.3:

- IPv4 IPFIX flow template ID—256
- IPv6 IPFIX flow template ID—257
- VPLS IPFIX flow template ID—258
- MPLS IPFIX flow template ID—259

The following are the default values of template IDs for version 9 flows for the different protocols or address families, starting with Junos OS Release 14.1:

- IPv4 version 9 flow template ID—320
- IPv6 version 9 flow template ID—321
- VPLS version 9 flow template ID—322
- MPLS version 9 flow template ID—323

The following are the default values of template IDs for IPFIX flows for the different protocols or address families, until Junos OS Release 13.3:

- IPv4 IPFIX flow options template ID—512
- IPv6 IPFIX flow options template ID—513
- VPLS IPFIX flow options template ID—514
- MPLS IPFIX flow options template ID—515



The following are the default values of template IDs for version 9 flows for the different protocols or address families, starting with Junos OS Release 14.1:

- IPv4 version 9 flow options template ID—576
- IPv6 version 9 flow options template ID—577
- VPLS version 9 flow options template ID—578
- MPLS version 9 flow options template ID—579

[Table 33 on page 923](#) describes the values of data template and option template IDs for different protocols with default and configured values for IPFIX flows.

**Table 43: Values of Template and Option Template IDs for IPFIX Flows**

Family	Configured Value	Data Template	Option Template
IPv4	None	256	576
IPv4	1024-65535	1024-65535	1024-65535
IPv6	None	257	577
IPv6	1024-65535	1024-65535	1024-65535
VPLS	None	258	578
VPLS	1024-65535	1024-65535	1024-65535
MPLS	None	259	579
MPLS	1024-65535	1024-65535	1024-65535

[Table 34 on page 923](#) describes the values of data template and option template IDs for different protocols with default and configured values for version 0 flows.

**Table 44: Values of Template and Option Template IDs for Version 9 Flows**

Family	Configured Value	Data Template	Option Template
IPv4	None	320	576
IPv4	1024-65535	1024-65535	1024-65535
IPv6	None	321	577
IPv6	1024-65535	1024-65535	1024-65535
VPLS	None	322	578
VPLS	1024-65535	1024-65535	1024-65535

**Table 44: Values of Template and Option Template IDs for Version 9 Flows (*continued*)**

Family	Configured Value	Data Template	Option Template
MPLS	None	323	579
MPLS	1024-65535	1024-65535	1024-65535

Table 33 on page 923 describes the values of data template and option template IDs for different protocols with default and configured values for IPFIX flows.

**Table 45: Values of Template and Option Template IDs for IPFIX Flows**

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id Conf val rsvd lproto slot LUInst PFEInst xxxx xxxx xxxx lxxx xxxx xxxx xxxx xxxx
None	IPv4 (0)	1	1	0	0000 0000 0000 1000 0000 0001 0000 0001 0x00080101
None	IPv6 (1)	1	1	0	0000 0000 0000 1001 0000 0001 0000 0001 0x00090101
None	VPLS (2)	1	1	0	0000 0000 0000 1010 0000 0001 0000 0001 0x000A0101
None	MPLS (3)	1	1	0	0000 0000 0000 1011 0000 0001 0000 0001 0x000B0101
4	IPv4 (0)	1	1	0	0000 0100 0000 1000 0000 0001 0000 0001 0x04080101

Table 45: Values of Template and Option Template IDs for IPFIX Flows (*continued*)

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id Conf val rsvd 1proto slot LUInst PFEInst  xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
190	IPv4 (0)	1	1	0	1101 1110 0000 1000 0000 0001 0000 0001 0xBE080101
4	IPv4 (0)	2	1	1	0000 0100 0000 1000 0000 0010 0001 0001 0x04080211
4	IPv6 (1)	1	1	0	0000 0100 0000 1001 0000 0001 0001 0000 0x04090110
190	IPv6 (1)	1	1	0	1101 1110 0000 1001 0000 0001 0001 0000 0xBE090110
4	VPLS (2)	2	2	0	0000 0100 0000 1010 0000 0010 0010 0000 0x040A0220
10	IPv4 (0)	28	2	1	0000 1010 0000 1000 0001 1100 0010 0001 0x0A081C21

**Related Documentation**

- [Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows on page 918](#)

## Directing Replicated Flows to Multiple Flow Servers

You can configure replication of the sampled flow records for use by multiple flow servers. You can use either sampling based on the Routing Engine, using cflowd version 5 or version 8, or sampling based on the services PIC, using flow aggregation version 9, as described in the following sections:

- [Directing Replicated Routing Engine–Based Sampling Flows to Multiple Servers on page 1106](#)
- [Directing Replicated Version 9 Flow Aggregates to Multiple Servers on page 1107](#)

### Directing Replicated Routing Engine–Based Sampling Flows to Multiple Servers

Routing Engine–based sampling supports up to eight flow servers for both cflowd version 5 and version 8 configurations. The total number of servers is limited to eight regardless of how many are configured for cflowd v5 or v8.

When you configure cflowd-based sampling, the export packets are replicated to all flow servers configured to receive them. If two servers are configured to receive v5 records, both the servers will receive records for a specified flow.



**NOTE:** With Routing Engine–based sampling, if multiple flow servers are configured with version 8 export format, all of them must use the same aggregation type. For example, all servers receiving version 8 export could be configured for source-destination aggregation type.

The following configuration example allows replication of export packets to two flow servers.

```
forwarding-options {
 sampling {
 instance inst1 {
 input {
 rate 1;
 }
 family inet;
 output {
 flow-server 10.10.3.2 {
 port 2055;
 version 5;
 source-address 192.168.164.119;
 }
 flow-server 172.17.20.62 {
 port 2055;
 version 5;
 source-address 192.168.164.119;
 }
 }
 }
 }
}
```

```
 }
 }
```

## Directing Replicated Version 9 Flow Aggregates to Multiple Servers

The export packets generated for a template are replicated to all the flow servers that are configured to receive information for that template. The maximum number of servers supported is eight.

This also implies that periodic updates required by version 9 (RFC 3954) are sent to each configured collector. The following updates are sent periodically as part of this requirement:

- Options data
- Template definition

The refresh period for options data and template definition is configured on a per-template basis at the **[edit services flow-monitoring]** hierarchy level.

The following configuration example allows replication of version 9 export packets to two flow servers.

```
forwarding-options {
 sampling {
 instance inst1 {
 input {
 rate 1;
 }
 family inet;
 output {
 flow-server 10.10.3.2 {
 port 2055;
 version9 {
 template {
 ipv4;
 }
 }
 }
 flow-server 172.17.20.62 {
 port 2055;
 version9 {
 template {
 ipv4;
 }
 }
 }
 }
 flow-inactive-timeout 30;
 flow-active-timeout 60;
 interface sp-4/0/0 {
 source-address 10.10.3.4;
 }
 }
 }
}
```

```
 }
 }
```

#### Related Documentation

- [Active Flow Monitoring Overview on page 815](#)
- [Configuring Flow Monitoring on page 818](#)
- [Configuring Services Interface Redundancy with Flow Monitoring on page 826](#)
- [Example: Configuring Active Monitoring on Logical Systems on page 823](#)

## Logging cflowd Flows Before Export

To collect the cflowd flows in a log file before they are exported, include the **local-dump** statement at the **[edit forwarding-options sampling output flow-server *hostname*]** hierarchy level:

```
[edit forwarding-options sampling output flow-server hostname]
local-dump;
```

By default, the flows are collected in **/var/log/sampled**; to change the filename, include the **filename** statement at the **[edit forwarding-options sampling traceoptions]** hierarchy level. For more information about changing the filename, see [“Configuring Traffic Sampling Output” on page 876](#).



**NOTE:** Because the **local-dump** statement adds extra overhead, you should use it only while debugging cflowd problems, not during normal operation.

The following is an example of the flow information. The AS number exported is the origin AS number. All flows that belong under a cflowd header are dumped, followed by the header itself:

```
Jun 27 18:35:43 v5 flow entry
Jun 27 18:35:43 Src addr: 192.53.127.1
Jun 27 18:35:43 Dst addr: 192.6.255.15
Jun 27 18:35:43 Nhop addr: 192.6.255.240
Jun 27 18:35:43 Input interface: 5
Jun 27 18:35:43 Output interface: 3
Jun 27 18:35:43 Pkts in flow: 15
Jun 27 18:35:43 Bytes in flow: 600
Jun 27 18:35:43 Start time of flow: 7230
Jun 27 18:35:43 End time of flow: 7271
Jun 27 18:35:43 Src port: 26629
Jun 27 18:35:43 Dst port: 179
Jun 27 18:35:43 TCP flags: 0x10
Jun 27 18:35:43 IP proto num: 6
Jun 27 18:35:43 TOS: 0xc0
Jun 27 18:35:43 Src AS: 7018
Jun 27 18:35:43 Dst AS: 11111
Jun 27 18:35:43 Src netmask len: 16
Jun 27 18:35:43 Dst netmask len: 0
```

[... 41 more version 5 flow entries; then the following header:]

```
Jun 27 18:35:43 cflowd header:
Jun 27 18:35:43 Num-records: 42
Jun 27 18:35:43 Version: 5
Jun 27 18:35:43 low seq num: 118
Jun 27 18:35:43 Engine id: 0
Jun 27 18:35:43 Engine type: 3
```

**Related  
Documentation**

- [Active Flow Monitoring Overview on page 815](#)
- [Configuring Flow Monitoring on page 818](#)
- [Directing Replicated Flows to Multiple Flow Servers on page 926](#)
- [Configuring Services Interface Redundancy with Flow Monitoring on page 826](#)
- [Example: Configuring Active Monitoring on Logical Systems on page 823](#)





# Sending Packets for Analysis Using Port Mirroring

- [Understanding Port Mirroring on page 1111](#)
- [Configuring Port Mirroring on page 1111](#)
- [Example: Multiple Port Mirroring with Next-Hop Groups Configuration on page 1128](#)

## Understanding Port Mirroring

---

On routers containing an Internet Processor II application-specific integrated circuit (ASIC) or T Series Internet Processor, you can send a copy of an IP version 4 (IPv4) or IP version 6 (IPv6) packet from the router to an external host address or a packet analyzer for analysis. This is known as *port mirroring*.

Port mirroring is different from traffic sampling. In traffic sampling, a sampling key based on the IPv4 header is sent to the Routing Engine. There, the key can be placed in a file, or cflowd packets based on the key can be sent to a cflowd server. In port mirroring, the entire packet is copied and sent out through a next-hop interface.

You can configure simultaneous use of sampling and port mirroring, and set an independent sampling rate and run-length for port-mirrored packets. However, if a packet is selected for both sampling and port mirroring, only one action can be performed and port mirroring takes precedence. For example, if you configure an interface to sample every packet input to the interface and a filter also selects the packet to be port mirrored to another interface, only the port mirroring would take effect. All other packets not matching the explicit filter port-mirroring criteria continue to be sampled when forwarded to their final destination.

### Related Documentation

- [Configuring Port Mirroring on page 931](#)
- [Example: Multiple Port Mirroring with Next-Hop Groups Configuration on page 949](#)

## Configuring Port Mirroring

---

To prepare traffic for port mirroring, include the **filter** statement at the **[edit firewall family inet]** hierarchy level:

```
filter filter-name;
```

This filter at the **[edit firewall family (inet | inet6)]** hierarchy level selects traffic to be port-mirrored:

```
filter filter-name {
 term term-name {
 then {
 port-mirror;
 accept;
 }
 }
}
```

To configure port mirroring on a logical interface, configure the following statements at the **[edit forwarding-options port-mirroring]** hierarchy level:

```
[edit forwarding-options port-mirroring family inet]
input {
 maximum-packet-length bytes;
 rate rate;
 run-length number;
}
family (inet|inet6) {
 output {
 interface interface-name {
 next-hop address;
 }
 no-filter-check;
 }
}
```

or

```
[edit forwarding-options port-mirroring]
input {
 maximum-packet-length bytes;
 rate rate;
 run-length number;
}
family inet6 {
 output {
 next-hop-group group-name {
 group-type inet6;
 interface interface-name {
 next-hop ipv6-address;
 }
 }
 next-hop-subgroup group-name {
 interface interface-name {
 next-hop ipv6-address;
 }
 }
 }
}
```



**NOTE:** The input statement can also be configured at the `[edit forwarding-options port-mirroring]` hierarchy level. This is only maintained for backward compatibility. However, the configuration of the output statement is deprecated at the `[edit forwarding-options port-mirroring]` hierarchy level.

Specify the port-mirroring destination by including the `next-hop` statement at the `[edit forwarding-options port-mirroring output interface interface-name]` hierarchy level:

```
next-hop address;
```



**NOTE:** For IPv4 port mirroring to reach a next-hop destination, you must manually include a static Address Resolution Protocol (ARP) entry in the router configuration.

You can also specify the port-mirroring destination by including the `next-hop-group` statement at the `[edit forwarding-options port-mirroring family inet6 output]` hierarchy level:

```
next-hop-group group-name {
 group-type inet6;
 interface interface-name {
 next-hop ipv6-address;
 }
 next-hop-subgroup group-name {
 interface interface-name {
 next-hop ipv6-address;
 }
 }
}
```

The `no-filter-check` statement is required when you send port-mirrored traffic to a Tunnel PIC that has a filter applied to it. en

The interface used to send the packets to the analyzer is the output interface configured above at the `[edit forwarding-options port-mirroring family (inet | inet6) output]` hierarchy level. You can use any physical interface type, including generic routing encapsulation (GRE) tunnel interfaces. The next-hop address specifies the destination address; this statement is mandatory for non point-to-point interfaces, such as Ethernet interfaces.

To configure the sampling rate or duration, include the `rate` or `run-length` statement at the `[edit forwarding-options port-mirroring input]` hierarchy level.

You can trace port-mirroring operations the same way you trace sampling operations. For more information, see [“Tracing Traffic Sampling Operations” on page 878](#).

For more information about port mirroring, see the following sections:

- [Configuring Tunnels on page 1114](#)
- [Port Mirroring with Next-Hop Groups on page 1116](#)

- [Configuring Inline Port Mirroring on page 1117](#)
- [Filter-Based Forwarding with Multiple Monitoring Interfaces on page 1118](#)
- [Restrictions on page 1118](#)
- [Configuring Port Mirroring on Services Interfaces on page 1119](#)
- [Examples: Configuring Port Mirroring on page 1120](#)

## Configuring Tunnels

In typical applications, you send the sampled packets to an analyzer or a workstation for analysis, rather than another router. If you must send this traffic over a network, you should use tunnels. For more information about tunnel interfaces, see *Tunnel Properties*.

The MX Series routers support Dense Port Concentrators (DPCs) with built-in Ethernet ports, which do not support Tunnel Services PICs. To create tunnel interfaces on an MX Series router with a DPC, you configure the DPC and the corresponding Packet Forwarding Engine to use for tunneling services at the **[edit chassis]** hierarchy level. You also configure the amount of bandwidth reserved for tunnel services. The Junos OS creates tunnel interfaces on the Packet Forwarding Engine.

To create tunnel interfaces on MX Series routers, include the following statements at the **[edit chassis]** hierarchy level:

```
[edit chassis]
fpc slot-number {
 pic number {
 tunnel-services {
 bandwidth bandwidth-value;
 }
 }
}
```

Include the **fpc slot-number** statement to specify the slot number of the DPC. If two SCBs are installed, the range is 0 through 1. If three SCBs are installed, the range is 0 through 5 and 7 through 11.

Include the **pic number** statement to specify the number of the Packet Forwarding Engine on the DPC. The range is 0 through 3.

You can also specify the amount of bandwidth to allocate for tunnel traffic on each Packet Forwarding Engine by including the **bandwidth bandwidth-value** statement at the **[edit chassis fpc slot-number pic pic-number]** hierarchy level:

- **1g** indicates that 1 Gbps of bandwidth is reserved for tunnel traffic. Configure this option only for a Packet Forwarding Engine on a Gigabit Ethernet 40-port DPC.
- **10g** indicates that 10 Gbps of bandwidth is reserved for tunnel traffic. Configure this option only for a Packet Forwarding Engine on a 10-Gigabit Ethernet 4-port DPC.
- **20g** or **40g**—Configure 20 gigabits per second or 40 gigabits per second only on an MX Series router with the MPC3E and the 100-Gigabit CFP MIC.

If you specify a bandwidth that is not compatible with the type of DPC and Packet Forwarding Engine, tunnel services are not activated. For example, you cannot specify a

bandwidth of 1 Gbps for a Packet Forwarding Engine on a 10-Gigabit Ethernet 4-port DPC.

When you configure tunnel interfaces on the Packet Forwarding Engine of a 10-Gigabit Ethernet 4-port DPC, the Ethernet interfaces for that port are removed from service and are no longer visible in the command-line interface (CLI). The Packet Forwarding Engine of a 10-Gigabit Ethernet 4-port DPC supports either tunnel interfaces or Ethernet interfaces, but not both. Each port on the 10-Gigabit Ethernet 4-port DPC includes two LEDs, one for tunnel services and one for Ethernet services, to indicate which type of service is being used. On the Gigabit Ethernet 40-port DPC, you can configure both tunnel and Ethernet interfaces at the same time.

If your router is equipped with a Tunnel PIC, you can forward duplicate packets to multiple interfaces by configuring a next-hop group. To configure a next-hop group, include the **next-hop-group** statement at the **[edit forwarding-options]** hierarchy level:

```
[edit forwarding-options]
next-hop-group group-names {
 interface interface-name {
 next-hop address;
 }
}
```

The **interface** statement specifies the interface that sends out sampled information. The **next-hop** statement specifies the next-hop addresses to which to send the sampled information.

For IPv6 port mirroring to reach next-hop destination, you can configure a **next-hop-group** statement at the **[edit forwarding-options port-mirroring family inet6 output]** hierarchy level:

```
next-hop-group group-name {
 group-type inet6;
 interface interface-name {
 next-hop ipv6-address;
 }
 next-hop-subgroup group-name {
 interface interface-name {
 next-hop ipv6-address;
 }
 }
}
```

Next-hop groups have the following restrictions:

- Next-hop groups are supported for inet, inet6, and bridge family.
- Next-hop groups are supported on M Series and MX Series routers.
- Next-hop groups or next-hop subgroups support up to 16 next-hop addresses.
- Up to 30 next-hop groups are supported.
- Each next-hop group is expected to have at least two next-hop addresses.
- Each next-hop subgroup supports up to 16 next-hop groups.

## Port Mirroring with Next-Hop Groups

You can configure next-hop groups for M Series, MX Series, and TX Series routers using either IP addresses or Layer 2 addresses for the next hops. Use the **group-type [ inet | inet6 | layer-2 ]** statement at **[edit forwarding-options next-hop-group next-hop-group-name]** hierarchy level to establish the next-hop groups. You can reference more than one port mirroring instance in a filter on MX Series routers. Use the **port-mirror-instance instance-name** statement at the **[edit firewall family family-name filter filter-name term term-name]** hierarchy level to refer to one of several port mirroring instances. For more information about this configuration, see the *Layer 2 Port Mirroring Feature Guide for Routing Devices*.



**NOTE:** On MX Series routers with MPCs, port mirroring instances can only be bound to the FPC level and not up to the PIC level. For MX series routers with a DPC card, both levels are supported.

On M Series, MX Series, and T Series routers only, you can configure port mirroring using next-hop groups, also known as *multipacket port mirroring*, without the presence of a Tunnel PIC. To configure this functionality, include the **next-hop-group** statement at the **[edit forwarding-options port-mirror family [inet | inet6] output]** or **[edit forwarding-options port-mirror instance instance-name family inet output]** hierarchy level:

```
[edit forwarding-options]
port-mirror {
 family inet {
 output {
 next-hop-group group-name {
 interface interface-name {
 next-hop address;
 }
 }
 }
 }
}
or
[edit forwarding-options]
port-mirror {
 family inet6 {
 output {
 next-hop-group group-name {
 group-type inet6;
 interface interface-name {
 next-hop ipv6-address;
 }
 }
 next-hop-subgroup group-name {
 interface interface-name {
 next-hop ipv6-address;
 }
 }
 }
 }
}
```

```

 }
 }
}
or
[edit forwarding-options]
port-mirror {
 instance instance-name {
 family (inet | vpls) {
 output {
 next-hop-group group-name;
 }
 }
 }
}

```

You define the next-hop group by including the **next-hop-group** statement at the **[edit forwarding-options]** hierarchy level. For an example, see [“Examples: Configuring Port Mirroring” on page 940](#). This configuration is supported with IPv4 and IPv6 addresses.

You can disable this configuration by including a **disable** or **disable-all-instances** statement at the **[edit forwarding-options port-mirror]** hierarchy level or by including a **disable** statement at the **[edit forwarding-options port-mirror instance *instance-name*]** hierarchy level. You can display the settings and network status by issuing the **show forwarding-options next-hop-group** and **show forwarding-options port-mirroring** operational commands.



**NOTE:** If you try to bind any derived instance to the FPC, a commit error will occur.

## Configuring Inline Port Mirroring

Inline port mirroring provides you with the ability to specify instances that are not bound to the flexible PIC concentrator (FPC) in the firewall filter’s **then port-mirror-instance** action. This way, you are not limited to only two port-mirror instances per FPC. Inline port mirroring decouples the port-mirror destination from the input parameters like **rate**. While the input parameters are programmed in the switch interface board, the next-hop destination of the mirrored packet is available in the packet itself. Inline port mirroring is supported only on MX Series routers with MPCs.

Using inline port mirroring, a port-mirror instance will have an option to inherit input parameters from another instance that specifies it, as shown in the following CLI configuration example:

```

instance pm2 {
 + input-parameters-instance pm1;
 family inet {
 output {
 interface ge-1/2/3.0 {
 next-hop 50.0.0.3;
 }
 }
 }
}

```

```
 }
 }
}
```

Multiple levels of inheritance are not allowed. One instance can be referred by multiple instances. An instance can refer to another instance that is defined before it. Forward references are not allowed and an instance cannot refer to itself, doing so will cause an error during configuration parsing.

The user can specify an instance that is not bound to the FPC in the firewall filter. The specified filter should inherit one of the two instances that have been bound to the FPC. If it does not, the packet is not marked for port-mirroring. If it does, then the packet will be sampled using the input parameters specified by the referred instance but the copy will be sent to the its own destination.

## Filter-Based Forwarding with Multiple Monitoring Interfaces

If port-mirrored packets are to be distributed to multiple monitoring or collection interfaces based on patterns in packet headers, it is helpful to configure a filter-based forwarding (FBF) filter on the port-mirroring egress interface.

When an FBF filter is installed as an output filter, a packet that is forwarded to the filter has already undergone at least one route lookup. After the packet is classified at the egress interface by the FBF filter, it is redirected to another routing table for additional route lookup. Obviously, the route lookup in the latter routing table (designated by an FBF routing instance) must result in a different next hop from those from the previous tables the packet has passed through, to avoid packet looping inside the Packet Forwarding Engine.

For more information about FBF configuration, see the *Junos OS Routing Protocols Library for Routing Devices*. For an example of FBF applied to an output interface, see [“Examples: Configuring Port Mirroring” on page 940](#).

## Restrictions

The following restrictions apply to port-mirroring configurations:

- The interface you configure for port mirroring should not participate in any kind of routing activity.
- The destination address you specify should not have a route to the ultimate traffic destination. For example, if the sampled IPv4 packets have a destination address of **10.68.9.10** and the port-mirrored traffic is sent to **10.68.20.15** for analysis, the device associated with the latter address should not know a route to **10.68.9.10**. Also, it should not send the sampled packets back to the source address.
- IPv4 and IPv6 traffic is supported. For IPv6 port mirroring, you must configure the next-hop router with an IPv6 neighbor before mirroring the traffic, similar to an ARP request for IPv4 traffic. All the restrictions applied to IPv4 configurations should also apply to IPv6.
- On M120 and M320 routers, multiple next-hop mirroring is not supported.



- Because M320 routers do not support multiple bindings of port-mirror instances per FPC, the **port-mirror-instance** action is not supported in firewall filters for these routers.
- Port mirroring in the ingress and egress direction is not supported for link services IQ (lsq-) interfaces.
- On M Series routers other than the M120 and M320 routers, only one family protocol (either IPv4 or IPv6) is supported at a time.
- Port mirroring supports up to 16 next hops.
- Only transit data is supported.
- You can configure multiple port-mirroring interfaces per router.
- On routers containing an Internet Processor II application-specific integrated circuit (ASIC), you must include a firewall filter with both the **accept** action and the **port-mirror** action modifier on the inbound interface. Do not include the **discard** action, or port mirroring will not work.
- If the port-mirroring interface is a non-point-to-point interface, you must include an IP address under the **port-mirroring** statement to identify the other end of the link. This IP address must be reachable for you to see the sampled traffic. If the port-mirroring interface is an Ethernet interface, the router should have an Address Resolution Protocol (ARP) entry for it. The following sample configuration sets up a static ARP entry.
- You do not need to configure firewall filters on both inbound and outbound interfaces, but at least one is necessary on the inbound interface to provide the copies of the packets to send to an analyzer.
- Inline port mirroring is supported only on MX Series routers with MPCs.
- Configuration for both port mirroring and traffic sampling are handled by the same daemon, so in order to view a trace log file for port mirroring, you must configure the **traceoptions** option under traffic sampling.

## Configuring Port Mirroring on Services Interfaces

A special situation arises when you configure unit **0** of a services interface (AS or Multiservices PIC) to be the port-mirroring logical interface, as in the following example:

```
[edit forwarding-options]
port-mirroring {
 input {
 rate 1;
 }
 family inet {
 output {
 interface sp-1/0/0.0;
 }
 }
}
```

Since any traffic directed to unit **0** on a services interface is targeted for monitoring (cflowd packets are generated for it), the sample port-mirroring configuration indicates

that the customer would like to have cflowd records generated for the port-mirrored traffic.

However, generation of cflowd records requires the following additional configuration; if it is missing, the port-mirrored traffic is simply dropped by the services interface without generating any cflowd packets.

```
[edit forwarding-options]
sampling {
 instance instance1 { # named instances of sampling parameters
 input {
 rate 1;
 }
 family inet {
 output {
 flow-server 172.16.28.65 {
 port 1230;
 }
 }
 interface sp-1/0/0 { # If the port-mirrored traffic requires monitoring, this
 # interface must be same as that specified in the
 # port-mirroring configuration.
 source-address 3.1.2.3;
 }
 }
 }
}
```



**NOTE:** Another way to configure sp-1/0/0 to generate cflowd records is to use only the sampling configuration, but include a firewall filter `sample` action instead of a port-mirror action.

---

## Examples: Configuring Port Mirroring

The following example sends port-mirrored traffic to multiple cflowd servers or packet analyzers:

```
[edit interfaces]
ge-1/0/0 { # This is the input interface where packets enter the router.
 unit 0 {
 family inet {
 filter {
 input mirror_pkts; # Here is where you apply the first filter.
 }
 address 10.11.0.1/24;
 }
 }
}
ge-1/1/0 { # This is an exit interface for HTTP packets.
 unit 0 {
 family inet {
 address 10.12.0.1/24;
 }
 }
}
```

```

 }
 }
 ge-1/2/0 { # This is an exit interface for HTTP packets.
 unit 0 {
 family inet {
 address 10.13.0.1/24;
 }
 }
 }
 so-0/3/0 { # This is an exit interface for FTP packets.
 unit 0 {
 family inet {
 address 10.1.1.1/30;
 }
 }
 }
 so-4/3/0 { # This is an exit interface for FTP packets.
 unit 0 {
 family inet {
 address 10.2.2.2/30;
 }
 }
 }
 so-7/0/0 { # This is an exit interface for all remaining packets.
 unit 0 {
 family inet {
 address 10.5.5.5/30;
 }
 }
 }
 so-7/0/1 { # This is an exit interface for all remaining packets.
 unit 0 {
 family inet {
 address 10.6.6.6/30;
 }
 }
 }
 vt-3/3/0 { # The tunnel interface is where you send the port mirrored traffic.
 unit 0 {
 family inet;
 }
 unit 1 {
 family inet {
 filter {
 input collect_pkts; # This is where you apply the second firewall filter.
 }
 }
 }
 }
}
[edit forwarding-options]
port-mirroring { # This is required when you configure next-hop groups.
 input {
 rate 1; # This rate port mirrors one packet for every one received (1:1 = all
 # packets).
 }
 family inet {

```

```
 output { # This sends traffic to a tunnel interface to prepare for multiport mirroring.
 interface vt-3/3/0.1;
 no-filter-check;
 }
 }
}
next-hop-group ftp-traffic { # Point-to-point interfaces require you to specify the interface
 # name only.
 interface so-4/3/0.0;
 interface so-0/3/0.0;
}
next-hop-group http-traffic { # You need to configure a next hop for multipoint interfaces
 # (Ethernet).
 interface ge-1/1/0.0 {
 next-hop 10.12.0.2;
 }
 interface ge-1/2/0.0 {
 next-hop 10.13.0.2;
 }
}
next-hop-group default-collect {
 interface so-7/0/0.0;
 interface so-7/0/1.0;
}
[edit firewall]
family inet {
 filter mirror_pkts { # Apply this filter to the input interface.
 term catch_all {
 then {
 count input_mirror_pkts;
 port-mirror; # This action sends traffic to be copied and port mirrored.
 accept;
 }
 }
 }
 filter collect_pkts { # Apply this filter to the tunnel interface.
 term ftp-term { # This term sends FTP traffic to an FTP next-hop group.
 from {
 protocol ftp;
 }
 then next-hop-group ftp-traffic;
 }
 term http-term { # This term sends HTTP traffic to an HTTP next-hop group.
 from {
 protocol http;
 }
 then next-hop-group http-traffic;
 }
 term default { # This term sends all remaining traffic to a final next-hop group.
 then next-hop-group default-collectors;
 }
 }
}
```

The following example demonstrates configuration of filter-based forwarding at the output interface. In this example, the packet flow follows this path:

1. A packet arrives at interface **fe-1/2/0.0** with source and destination addresses **10.50.200.1** and **10.50.100.1**, respectively.
2. The route lookup in routing table **inet.0** points to the egress interface **so-0/0/3.0**.
3. The output filter installed at **so-0/0/3.0** redirects the packet to routing table **fbf.inet.0**.
4. The packet matches the entry **10.50.100.0/25**, and finally leaves the router from interface **so-2/0/0.0**.

```
[edit interfaces]
so-0/0/3 {
 unit 0 {
 family inet {
 filter {
 output fbf;
 }
 address 10.50.10.2/25;
 }
 }
}
fe-1/2/0 {
 unit 0 {
 family inet {
 address 10.50.50.2/25;
 }
 }
}
so-2/0/0 {
 unit 0 {
 family inet {
 address 10.50.20.2/25;
 }
 }
}
[edit firewall]
filter fbf {
 term 0 {
 from {
 source-address {
 10.50.200.0/25;
 }
 }
 then routing-instance fbf;
 }
 term d {
 then count d;
 }
}
[edit routing-instances]
fbf {
 instance-type forwarding;
 routing-options {
```

```
 static {
 route 10.50.100.0/25 next-hop so-2/0/0.0;
 }
 }
}
[edit routing-options]
interface-routes {
 rib-group inet fbf-group;
}
static {
 route 10.50.100.0/25 next-hop 10.50.10.1;
}
rib-groups {
 fbf-group {
 import-rib [inet.0 fbf.inet.0];
 }
}
```

The following example shows configuration of port mirroring using next-hop groups or multipacket port mirroring:

```
forwarding-options {
 next-hop-group inet_nhg {
 group-type inet;
 interface ge-2/0/2.101 {
 next-hop 10.2.0.2;
 }
 interface ge-2/2/8.2 {
 next-hop 10.8.0.2;
 }
 }
 next-hop-group vpls_nhg {
 group-type layer-2;
 interface ge-2/0/1.100;
 interface ge-2/2/9.0;
 inactive: next-hop-subgroup vpls_subg {
 interface ge-2/0/1.101;
 interface ge-2/2/9.1;
 }
 }
 next-hop-group vpls_nhg_2 {
 group-type layer-2;
 interface ge-2/2/1.100;
 interface ge-2/3/9.0;
 }
}
port-mirror {
 disable-all-instances; /* Disable all port-mirroring instances */
 disable; /* Disable the global instance */
 input {
 rate 10; # start mirroring every 10th packet
 run-length 4; # mirror 4 additional packets
 }
 family inet {
 output {
 next-hop-group inet_nhg;
 }
 }
}
```

```

}
family inet6 {
 output {
 next-hop-group inet6_nhg6 {
 group-type inet6;
 interface ge-2/0/3.102 {
 next-hop 10::1:1:10;
 }
 interface ge-2/0/4.103 {
 next-hop 20::1:1:10;
 }
 next-hop-subgroup vpls_subg {
 interface ge-2/0/.101 {
 next-hop 3::1:1:1;
 }
 interface ge-2/2/9.1 {
 next-hop 4::1:1:1;
 }
 }
 }
 }
}
family vpls {
 output {
 next-hop-group vpls_nhg;
 }
}
instance {
 inst1 {
 disable; /* Disable this instance */
 input {
 rate 1;
 maximum-packet-length 200;
 }
 family inet {
 output {
 next-hop-group inet_nhg;
 }
 }
 family inet6 {
 output {
 next-hop-group inet6_nhg6;
 }
 }
 family vpls {
 output {
 next-hop-group vpls_nhg_2;
 }
 }
 }
}
}

```

The following example shows configuration of port mirroring using next-hop groups or multipacket port mirroring on a T Series router:

```
forwarding-options {
 next-hop-group inet_nhg {
 group-type inet;
 interface so-0/0/0.0; # There is no need for the nexthop address on T Series routers
 interface ge-2/0/2.0 {
 next-hop 1.2.3.4
 }
 }
 next-hop-subgroup sub_inet {
 interface so-1/2/0.0;
 interface ge-6/1/2.0 {
 next-hop 6.7.8.9;
 }
 }
 next-hop-group vpls_nhg_2 {
 group-type layer-2;
 interface ge-2/2/1.100;
 interface ge-2/3/9.0;
 }
}
port-mirroring {
 disable-all-instances; /*Disable all port-mirroring instances */
 disable; /* Disable the global instance */
 input {
 rate 10;
 run-length 4;
 }
 family inet {
 output {
 next-hop-group inet_nhg;
 }
 }
 family vpls {
 output {
 next-hop-group vpls_nhg;
 }
 }
 instance {
 inst1 {
 disable; /* Disable this instance */
 input {
 rate 1;
 maximum-packet-length 200;
 }
 family inet {
 output {
 next-hop-group inet_nhg;
 }
 }
 family vpls {
 output {
 next-hop-group vpls_nhg_2;
 }
 }
 }
 }
}
```



```
}

```

The following example shows configuration of inline port mirroring using PM1, PM2, and PM3 as our port mirror instances.

```
instance {
 pm1 {
 input {
 rate 3;
 }
 family inet {
 output {
 interface ge-1/2/2.0 {
 next-hop 40.0.0.2;
 }
 }
 }
 }
 pm2 {
 input-parameters-instance pm1;
 family inet {
 output {
 interface ge-1/2/3.0 {
 next-hop 50.0.0.3;
 }
 }
 }
 }
 pm3 {
 input {
 rate 3;
 }
 family inet6 {
 output {
 interface ge-1/2/3.0 {
 next-hop 5::5:5:1;
 }
 }
 }
 }
}
firewall {
 filter pm_filter {
 term t1 {
 then port-mirror-instance pm2;
 }
 }
 filter nhg6_filter6 {
 term t6 {
 then next-hop-group inet6-nhg6;
 }
 }
}
chassis {
 fpc 1 {
 port-mirror-instance pm1;
 }
}
```

```
}
```

The packets will be sampled at a rate of 3, and the copy is sent to 50.0.0.3.

**Related  
Documentation**

- [Understanding Port Mirroring on page 931](#)
- [Example: Multiple Port Mirroring with Next-Hop Groups Configuration on page 949](#)

---

## Example: Multiple Port Mirroring with Next-Hop Groups Configuration

When you need to analyze traffic containing more than one packet type, or you wish to perform multiple types of analysis on a single type of traffic, you can implement multiple port mirroring and next-hop groups. You can make up to 16 copies of traffic per group and send the traffic to next-hop group members. A maximum of 30 groups can be configured on a router at any given time. The port-mirrored traffic can be sent to any interface, except aggregated SONET/SDH, aggregated Ethernet, loopback (**lo0**), or administrative (**fxp0**) interfaces. To send port-mirrored traffic to multiple flow servers or packet analyzers, you can use the **next-hop-group** statement at the **[edit forwarding-options]** hierarchy level.

Figure 42: Active Flow Monitoring—Multiple Port Mirroring with Next-Hop Groups Topology Diagram




Figure 35 on page 950 shows an example of how to configure multiple port mirroring with next-hop groups. All traffic enters the monitoring router at interface **ge-1/0/0**. A firewall filter counts and port-mirrors all incoming packets to a Tunnel Services PIC. A second filter is applied to the tunnel interface and splits the traffic into three categories: HTTP traffic, FTP traffic, and all other traffic. The three types of traffic are assigned to three separate next-hop groups. Each next-hop group contains a unique pair of exit interfaces that lead to different groups of packet analyzers and flow servers.



**NOTE:** Instances enabled to mirror packets to different destinations from the same PFE, also use different sampling parameters for each instance. When we configure Layer2 Port-mirroring with both global port-mirroring and instance based port-mirroring, PIC level instances will override FPC level and the FPC level will override the Global instance.

[edit]  
interfaces {

```
ge-1/0/0 { # This is the input interface where packets enter the router.
 unit 0 {
 family inet {
 filter {
 input mirror_pkts; # Here is where you apply the first filter.
 }
 address 10.11.1.1/24;
 }
 }
}
ge-1/1/0 { # This is an exit interface for HTTP packets.
 unit 0 {
 family inet {
 address 10.12.1.1/24;
 }
 }
}
ge-1/2/0 { # This is an exit interface for HTTP packets.
 unit 0 {
 family inet {
 address 10.13.1.1/24;
 }
 }
}
so-0/3/0 { # This is an exit interface for FTP packets.
 unit 0 {
 family inet {
 address 10.1.1.1/30;
 }
 }
}
so-4/3/0 { # This is an exit interface for FTP packets.
 unit 0 {
 family inet {
 address 10.2.2.1/30;
 }
 }
}
so-7/0/0 { # This is an exit interface for all remaining packets.
 unit 0 {
 family inet {
 address 10.5.5.1/30;
 }
 }
}
so-7/0/1 { # This is an exit interface for all remaining packets.
 unit 0 {
 family inet {
 address 10.6.6.1/30;
 }
 }
}
vt-3/3/0 { # The tunnel interface is where you send the port-mirrored traffic.
 unit 0 {
 family inet;
 }
}
```

```

unit 1 {
 family inet {
 filter {
 input collect_pkts; # This is where you apply the second firewall filter.
 }
 }
}
}
forwarding-options {
 port-mirroring { # This is required when you configure next-hop groups.
 family inet {
 input {
 rate 1; # This port-mirrors all packets (one copy for every packet received).
 }
 output { # Sends traffic to a tunnel interface to enable multipoint mirroring.
 interface vt-3/3/0.1;
 no-filter-check;
 }
 }
 }
 next-hop-group ftp-traffic { # Point-to-point interfaces require you to specify the
 interface so-4/3/0.0; # interface name.
 interface so-0/3/0.0;
 }
 next-hop-group http-traffic { # Configure a next hop for all multipoint interfaces.
 interface ge-1/1/0.0 {
 next-hop 10.12.1.2;
 }
 interface ge-1/2/0.0 {
 next-hop 10.13.1.2;
 }
 }
 next-hop-group default-collect {
 interface so-7/0/0.0;
 interface so-7/0/1.0;
 }
}
}
firewall {
 family inet {
 filter mirror_pkts { # Apply this filter to the input interface.
 term catch_all {
 then {
 count input_mirror_pkts;
 port-mirror; # This action sends traffic to be copied and port-mirrored.
 }
 }
 }
 filter collect_pkts { # Apply this filter to the tunnel interface.
 term ftp-term { # This term sends FTP traffic to an FTP next-hop group.
 from {
 protocol ftp;
 }
 then next-hop-group ftp-traffic;
 }
 term http-term { # This term sends HTTP traffic to an HTTP next-hop group.

```

```
 from {
 protocol http;
 }
 then next-hop-group http-traffic;
 }
 term default { # This sends all remaining traffic to a final next-hop group.
 then next-hop-group default-collectors;
 }
}
}
```

- Related Documentation**
- [Understanding Port Mirroring on page 931](#)
  - [Configuring Port Mirroring on page 931](#)

## PART 19

# Real-Time Performance Monitoring and Video Monitoring Services

- [Monitoring Traffic Using Real-Time Performance Monitoring on page 1135](#)
- [Testing the Performance of Network Devices Using RFC 2544-Based Benchmarking on page 1161](#)
- [Tracking Streaming Media Traffic Using Inline Video Monitoring on page 1191](#)





## CHAPTER 71

# Monitoring Traffic Using Real-Time Performance Monitoring

- [Real-Time Performance Monitoring Services Overview on page 1135](#)
- [Configuring RPM Probes on page 1137](#)
- [Configuring RPM Receiver Servers on page 1141](#)
- [Limiting the Number of Concurrent RPM Probes on page 1142](#)
- [Configuring RPM Timestamping on page 1142](#)
- [Configuring TWAMP on page 1146](#)
- [Configuring BGP Neighbor Discovery Through RPM on page 1149](#)
- [Examples: Configuring BGP Neighbor Discovery Through RPM on page 1151](#)
- [Tracing RPM Operations on page 1153](#)
- [Examples: Configuring Real-Time Performance Monitoring on page 1155](#)
- [Enabling RPM for the Junos OS extension-provider package on page 1159](#)

## Real-Time Performance Monitoring Services Overview

---

Real-Time Performance Monitoring (RPM) enables you to configure active probes to track and monitor traffic. Probes collect packets per destination and per application, including PING Internet Control Message Protocol (ICMP) packets, User Datagram Protocol and Transmission Control Protocol (UDP/TCP) packets with user-configured ports, user-configured Differentiated Services code point (DSCP) type-of-service (ToS) packets, and Hypertext Transfer Protocol (HTTP) packets. RPM provides Management Information Base (MIB) support with extensions for RFC 2925, *Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations*.

You can also configure RPM services to determine automatically whether a path exists between a host router and its configured BGP neighbors. You can view the results of the discovery using an SNMP client. Results are stored in **pingResultsTable**, **jnxPingResultsTable**, **jnxPingProbeHistoryTable**, and **pingProbeHistoryTable**.

Probe configuration and probe results are supported by the command-line interface (CLI) and SNMP.

The following probe types are supported with DSCP marking:

- ICMP echo
- ICMP timestamp
- HTTP get (not available for BGP RPM services)
- UDP echo
- TCP connection
- UDP timestamp

With probes, you can monitor the following:

- Minimum round-trip time
- Maximum round-trip time
- Average round-trip time
- Standard deviation of the round-trip time
- Jitter of the round-trip time—The difference between the minimum and maximum round-trip time

One-way measurements for ICMP timestamp probes include the following:

- Minimum, maximum, standard deviation, and jitter measurements for egress and ingress times
- Number of probes sent
- Number of probe responses received
- Percentage of lost probes



**NOTE:** Timestamping is not supported on PTX Series Packet Transport Routers.

---

You can configure the following RPM thresholds:

- Round-trip time
- Ingress/egress delay
- Standard deviation
- Jitter
- Successive lost probes
- Total lost probes (per test)

Support is also implemented for user-configured CoS classifiers and for prioritization of RPM packets over regular data packets received on an input interface.

- Related Documentation**
- [Configuring BGP Neighbor Discovery Through RPM on page 971](#)
  - [\[edit services rpm\] Hierarchy Level on page 1631](#)
  - [Examples: Configuring BGP Neighbor Discovery Through RPM on page 973](#)

## Configuring RPM Probes

The owner name and test name identifiers of an RPM probe together represent a single RPM configuration instance. When you specify the test name, you also can configure the test parameters.

To configure the probe owner, test name, and test parameters, include the **probe** statement at the **[edit services rpm]** hierarchy level:

```
probe owner {
 test test-name {
 data-fill data;
 data-size size;
 destination-interface interface-name;
 destination-port port;
 dscp-code-point dscp-bits;
 hardware-timestamp;
 history-size size;
 moving-average-size number;
 one-way-hardware-timestamp;
 probe-count count;
 probe-interval seconds;
 probe-type type;
 routing-instance instance-name;
 source-address address;
 target (url url | address address);
 test-interval interval;
 thresholds thresholds;
 traps traps;
 }
}
```

Keep the following points in mind when you configure RPM clients and RPM servers:

- You cannot configure an RPM client that is PIC-based and an RPM server that is based on either the Packet Forwarding Engine or Routing Engine to receive the RPM probes.
- You cannot configure an RPM client that is Packet Forwarding Engine-based and an RPM server that receives the RPM probes to be on the PIC or Routing Engine.
- The RPM client and RPM server must be located on the same type of module. For example, if the RPM client is PIC-based, the RPM server must also be PIC-based, and if the RPM server is Packet Forwarding Engine-based, the RPM client must also be Packet Forwarding Engine-based.

- To specify a probe owner, include the **probe** statement at the **[edit services rpm]** hierarchy level. The probe owner identifier can be up to 32 characters in length.
- To specify a test name, include the **test** statement at the **[edit services rpm probe owner]** hierarchy level. The test name identifier can be up to 32 characters in length. A test represents the range of probes over which the standard deviation, average, and jitter are calculated.
- To specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes, include the **data-fill** statement at the **[edit services rpm probe owner]** hierarchy level. The value can be a hexadecimal value. The **data-fill** statement is not valid with the **http-get** or **http-metadata-get** probe types.
- To specify the size of the data portion of ICMP probes, include the **data-size** statement at the **[edit services rpm probe owner]** hierarchy level. The size can be from 0 through 65400 and the default size is 0. The **data-size** statement is not valid with the **http-get** or **http-metadata-get** probe types.



**NOTE:** If you configure the hardware timestamp feature (see [“Configuring RPM Timestamping” on page 964](#)):

- The **data-size** default value is 32 bytes and 32 is the minimum value for explicit configuration. The UDP timestamp probe type is an exception; it requires a minimum data size of 44 bytes.
  - The **data-size** must be at least 100 bytes smaller than the default MTU of the interface of the RPM client interface.
- 
- On M Series and T Series routers, you configure the **destination-interface** statement to enable hardware timestamping of RPM probe packets. You specify an **sp-** interface to have the AS or Multiservices PIC add the hardware timestamps; for more information, see [“Configuring RPM Timestamping” on page 964](#). You can also include the **one-way-hardware-timestamp** statement to enable one-way delay and jitter measurements.
  - To specify the User Datagram Protocol (UDP) port or Transmission Control Protocol (TCP) port to which the probe is sent, include the **destination-port** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. The **destination-port** statement is used only for the UDP and TCP probe types. The value can be 7 or from 49160 through 65535.

When you configure either **probe-type udp-ping** or **probe-type udp-ping-timestamp** along with hardware timestamping, the value for the **destination-port** can be only 7. A constraint check prevents you from configuring any other value for the destination port in this case. This constraint does not apply when you are using one-way hardware timestamping.

- To specify the value of the Differentiated Services (DiffServ) field within the IP header, include the **dscp-code-point** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. The DiffServ code point (DSCP) bits value can be set to a valid 6-bit pattern; for example, 001111. It also can be set using an alias configured at

the **[edit class-of-service code-point-aliases dscp]** hierarchy level. The default is 000000.

- To specify the number of stored history entries, include the **history-size** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. Specify a value from 0 to 512. The default is 50.
- To specify a number of samples for making statistical calculations, include the **moving-average-size** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. Specify a value from 0 through 255.
- To specify the number of probes within a test, include the **probe-count** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. Specify a value from 1 through 15.
- To specify the time to wait between sending packets, include the **probe-interval** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. Specify a value from 1 through 255 seconds.
- To specify the packet and protocol contents of the probe, include the **probe-type** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. The following probe types are supported:
  - **http-get**—Sends a Hypertext Transfer Protocol (HTTP) get request to a target URL.
  - **http-metadata-get**—Sends an HTTP get request for metadata to a target URL.
  - **icmp-ping**—Sends ICMP echo requests to a target address.
  - **icmp-ping-timestamp**—Sends ICMP timestamp requests to a target address.
  - **tcp-ping**—Sends TCP packets to a target.
  - **udp-ping**—Sends UDP packets to a target.
  - **udp-ping-timestamp**—Sends UDP timestamp requests to a target address.

The following probe types support hardware timestamping of probe packets: **icmp-ping**, **icmp-ping-timestamp**, **udp-ping**, **udp-ping-timestamp**.



**NOTE:** Some probe types require additional parameters to be configured. For example, when you specify the **tcp-ping** or **udp-ping** option, you must configure the destination port using the **destination-port** statement. The **udp-ping-timestamp** option requires a minimum data size of 12; any smaller data size results in a commit error. The minimum data size for TCP probe packets is 1.

When you configure either **probe-type udp-ping** or **probe-type udp-ping-timestamp** along with the **one-way-hardware-timestamp** command, the value for the **destination-port** can be only 7. A constraint check prevents you for configuring any other value for the destination port in this case.

- To specify the routing instance used by ICMP probes, include the **routing-instance** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. The default routing instance is Internet routing table **inet.0**.
- To specify the source IP address used for ICMP probes, include the **source-address** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. If the source IP address is not one of the router's assigned addresses, the packet will use the outgoing interface's address as its source.
- To specify the destination address used for the probes, include the **target** statement at the **[edit services rpm probe owner test test-name]** hierarchy level.
  - For HTTP probe types, specify a fully formed URL that includes **http://** in the URL address.
  - For all other probe types, specify an IP version 4 (IPv4) address for the target host.
- To specify the time to wait between tests, include the **test-interval** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. Specify a value from 1 through 86400 seconds.



**NOTE:** Starting with Junos OS Release 15.1, the minimum period for which the RPM client waits between two tests is modified to be 1 second instead of 0 seconds. Also, if you do not configure the test interval, the default value is 1 second.

- To specify thresholds used for the probes, include the **thresholds** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. A system log message is generated when the configured threshold is exceeded. Likewise, an SNMP trap (if configured) is generated when a threshold is exceeded. The following options are supported:
  - **egress-time**—Measures maximum source-to-destination time per probe.
  - **ingress-time**—Measures maximum destination-to-source time per probe.
  - **jitter-egress**—Measures maximum source-to-destination jitter per test.
  - **jitter-ingress**—Measures maximum destination-to-source jitter per test.
  - **jitter-rtt**—Measures maximum jitter per test, from 0 through 60000000 microseconds.
  - **rtt**—Measures maximum round-trip time per probe, in microseconds.
  - **std-dev-egress**—Measures maximum source-to-destination standard deviation per test.
  - **std-dev-ingress**—Measures maximum destination-to-source standard deviation per test.
  - **std-dev-rtt**—Measures maximum standard deviation per test, in microseconds.

- **successive-loss**—Measures successive probe loss count, indicating probe failure.
- **total-loss**—Measures total probe loss count indicating test failure, from 0 through 15.
- Traps are sent if the configured threshold is met or exceeded. To set the trap bit to generate traps, include the **traps** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. The following options are supported:
  - **egress-jitter-exceeded**—Generates traps when the jitter in egress time threshold is met or exceeded.
  - **egress-std-dev-exceeded**—Generates traps when the egress time standard deviation threshold is met or exceeded.
  - **egress-time-exceeded**—Generates traps when the maximum egress time threshold is met or exceeded.
  - **ingress-jitter-exceeded**—Generates traps when the jitter in ingress time threshold is met or exceeded.
  - **ingress-std-dev-exceeded**—Generates traps when the ingress time standard deviation threshold is met or exceeded.
  - **ingress-time-exceeded**—Generates traps when the maximum ingress time threshold is met or exceeded.
  - **jitter-exceeded**—Generates traps when the jitter in round-trip time threshold is met or exceeded.
  - **probe-failure**—Generates traps for successive probe loss thresholds crossed.
  - **rtt-exceeded**—Generates traps when the maximum round-trip time threshold is met or exceeded.
  - **std-dev-exceeded**—Generates traps when the round-trip time standard deviation threshold is met or exceeded.
  - **test-completion**—Generates traps when a test is completed.
  - **test-failure**—Generates traps when the total probe loss threshold is met or exceeded.

**Related  
Documentation**

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [Examples: Configuring Real-Time Performance Monitoring on page 977](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)

---

## Configuring RPM Receiver Servers

The RPM TCP and UDP probes are proprietary to Juniper Networks and require a receiver to receive the probes. To configure a server to receive the probes, include the **probe-server** statement at the **[edit services rpm]** hierarchy level:

```
[edit services rpm]
probe-server {
```

```
tcp {
 destination-interface interface-name;
 port number;
}
udp {
 port number;
}
}
```

The port number specified for the UDP and TCP server can be 7 or from 49160 through 65535.



**NOTE:** The `destination-interface` statement is not supported on PTX Series Packet Transport Routers.

When you configure either `probe-type udp-ping` or `probe-type udp-ping-timestamp` along with the `one-way-hardware-timestamp` command, the value for the `destination-port` can be only 7. A constraint check prevents you for configuring any other value for the destination port in this case.

**Related  
Documentation**

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)
- [Examples: Configuring Real-Time Performance Monitoring on page 977](#)

---

## Limiting the Number of Concurrent RPM Probes

To configure the maximum number of concurrent probes allowed, include the `probe-limit` statement at the `[edit services rpm]` hierarchy level:

```
probe-limit limit;
```

Specify a limit from 1 through 500. The default maximum number is 100.

**Related  
Documentation**

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)
- [Examples: Configuring Real-Time Performance Monitoring on page 977](#)

---

## Configuring RPM Timestamping

To account for latency in the communication of probe messages, you can enable timestamping of the probe packets. You can timestamp the following RPM probe types: `icmp-ping`, `icmp-ping-timestamp`, `udp-ping`, and `udp-ping-timestamp`.



On M Series and T Series routers with an Adaptive Services (AS) or Multiservices PIC, and on EX Series switches with a Multiservices DPC, and on EX Series switches, you can enable hardware timestamping of RPM probe messages. The timestamp is applied on both the RPM client router (the router or switch that originates the RPM probes) and the RPM probe server and applies only to IPv4 traffic. It is supported on the following:

- Layer 2 services package on all Multiservices PICs and DPCs.
- Layer 3 service package on AS and Multiservices PICs and Multiservices DPCs.
- Extension-provider services package on M Series, MX Series, and T Series services PICs that support the Extension-Provider packages (In Junos OS releases earlier than 12.3, the extension-provider packages were variously referred to as Junos Services Framework (JSF), MP-SDK, and eJunos.)
- Layer 2, Layer 3, SDK Services, and PFE RPM timestamping interoperate with each other. Here, the RPM client can be on the Layer 3 **sp-** interface and the RPM server can be on an SDK Services package.



**NOTE:** Hardware timestamping is not supported on PTX Series Packet Transport Routers.

Two-way timestamping is available on **sp-** and **ms-** interfaces. To configure two-way timestamping on M Series and T Series routers, include the **destination-interface** statement at the **[edit services rpm probe probe-owner test test-name]** hierarchy level:

```
destination-interface sp-fpc/pic/port.logical-unit
destination-interface ms-fpc/pic/port.logical-unit
```

Specify the RPM client router and the RPM server router on the adaptive services logical interface or the multiservices interface by including the **rpm** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level:

```
rpm (client | server);
```

The logical interface must be dedicated to the RPM task. It requires configuration of the **family inet** statement and a **/32** address, as shown in the example. This configuration is also needed for other services such as NAT and stateful firewall. You cannot configure RPM service on **unit 0** because RPM requires a dedicated logical interface; the same unit cannot support both RPM and other services. Because active flow monitoring requires **unit 0**, but RPM can function on any logical interface, a constraint check prevents you from committing an RPM configuration there.



**NOTE:** If you configure RPM timestamping on an AS PIC, you cannot configure the **source-address** statement at the **[edit services rpm probe probe-name test test-name]** hierarchy level.

On MX Series routers, on M-320 routers using the Enhanced Queuing MPC, and on EX Series switches, you include the **hardware-timestamp** statement at the **[edit services rpm**

**probe *probe-name* test *test-name*]** hierarchy level to specify that the probes are to be timestamped in the Packet Forwarding Engine host processor:

**hardware-timestamp;**

On the client side, these probes are timestamped in the Packet Forwarding Engine host processor on the egress DPC on the MX or M-320 Series router or EX Series switch originating the RPM probes (RPM client). On the responder side (RPM server), the RPM probes to be timestamped are handled by the Packet Forwarding Engine host processor, which generates the response instead of the RPM process. The RPM probes are timestamped only on the router that originates them (RPM client). As a result, only round-trip time is measured for these probes.

When using the **hardware-timestamp**, the **data-size** value for the probe must be at least 100 bytes smaller than the default MTU of the interface of the RPM client interface (see [“Configuring RPM Probes” on page 959](#)). If hardware timestamping of RPM probe messages is enabled, the maximum data size that you can configure by using the data-size statement is limited to 1400.



**NOTE:** The Packet Forwarding Engine-based RPM feature does not support any stateful firewall configurations. If you need to combine RPM timestamping with a stateful firewall, you should use the interface-based RPM timestamping service described earlier in this section. Multiservices DPCs support stateful firewall processing as well as RPM timestamping.

---

To configure one-way timestamping, you must also include the **one-way-hardware-timestamp** statement at the **[edit services rpm probe *probe-owner* test *test-name*]** hierarchy level:

**one-way-hardware-timestamp;**



**NOTE:** If you configure RPM probes for a services interface (sp-), you need to announce local routes in a specific way for the following routing protocols:

- For OSPF, you can announce the local route by including the services interface in the OSPF area. To configure this setting, include the interface `sp-fpc/pic/port` statement at the [edit protocols ospf area *area-number*] hierarchy level.
- For BGP and IS-IS, you must export interface routes and create a policy that accepts the services interface local route. To export interface routes, include the point-to-point and lan statements at the [edit routing-options interface-routes family inet export] hierarchy level. To configure an export policy that accepts the services interface local route, include the protocol local, rib inet.0, and route-filter `sp-interface-ip-address/32` exact statements at the [edit policy-options policy-statement *policy-name* term *term-name* from] hierarchy level and the accept action at the [edit policy-options policy-statement *policy-name* term *term-name* then] hierarchy level. For the export policy to take effect, apply the policy to BGP or IS-IS with the export *policy-name* statement at the [edit protocols *protocol-name*] hierarchy level.

For more information about these configurations, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices* or the *Junos OS Routing Protocols Library for Routing Devices*.

Routing the probe packets through the adaptive services or Multiservices PIC also enables you to filter the probe packets to particular queues. The following example shows the RPM configuration and the filter that specifies queuing:

```
services rpm {
 probe p1 {
 test t1 {
 probe-type icmp-ping;
 target address 10.8.4.1;
 probe-count 10;
 probe-interval 10;
 test-interval 10;
 dscp-code-points af11;
 data-size 100;
 destination-interface sp-1/2/0.0;
 }
 }
}
firewall {
 filter f1 {
 term t1 {
 from {
 dscp af11;
 }
 then {
 forwarding-class assured-forwarding;
 }
 }
 }
}
```

```
 }
 }
}
interfaces sp-1/2/0 {
 unit 2 {
 rpm client;
 family inet {
 address 10.8.4.2/32;
 filter {
 input f1;
 }
 }
 }
}
interfaces sp-1/2/1 {
 unit 2 {
 rpm server;
 family inet {
 address 10.8.3.2/32;
 filter {
 input f1;
 }
 }
 }
}
```

For more information about firewall filters, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*; for more information about queuing, see the *Class of Service Feature Guide for Routing Devices*.

#### Related Documentation

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)
- [Examples: Configuring Real-Time Performance Monitoring on page 977](#)

---

## Configuring TWAMP

You can configure the Two-Way Active Measurement Protocol (TWAMP) on all M Series and T Series routers that support Multiservices PICs (running in either Layer 2 or Layer 3 mode), and on MX Series routers. Only the responder (server) side of TWAMP is supported.



**NOTE:** TWAMP is not supported on EX Series switches and PTX Series Packet Transport Routers.

For more information on TWAMP, see RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP)*.

To configure TWAMP properties, include the **twamp** statement at the [\[edit services rpm\]](#) hierarchy level:

```
[edit services rpm]
```

```

twamp {
 server {
 client-list list-name {
 [address address];
 }
 authentication-mode mode;
 max-connection-duration hours;
 maximum-connections count;
 maximum-connections-per-client count;
 maximum-sessions count;
 maximum-sessions-per-connection count;
 port number;
 routing-instance-list {
 instance-name {
 port number;
 }
 }
 server-inactivity-timeout minutes;
 }
}

```

The TWAMP configuration process includes the following tasks:

- [Configuring TWAMP Interfaces on page 1147](#)
- [Configuring TWAMP Servers on page 1147](#)

## Configuring TWAMP Interfaces

To specify the service PIC logical interface that provides the TWAMP service, include the **twamp-server** statement at the **[edit interfaces sp-fpc/pic/port unit logical-unit-number]** hierarchy level:

```
twamp-server;
```



**NOTE:** On MX Series routers that do not include a Multiservices DPC, you can configure the **twamp-server** statement on any interface (for example, ge-1/0/1.10). It is not necessary to configure this statement on a service interface (sp- or ms-) but you do need to include it in the configuration to activate the TWAMP reflector functionality.

## Configuring TWAMP Servers

You can specify a number of TWAMP server properties, some of which are optional, by including the **server** statement at the **[edit services rpm twamp]** hierarchy level:

```

[edit services rpm twamp]
server {
 client-list list-name {
 [address address];
 }
 authentication-mode mode;
 max-connection-duration hours;
 maximum-connections count;

```

```
maximum-connections-per-client count;
maximum-sessions count;
maximum-sessions-per-connection count;
port number;
routing-instance-list {
 instance-name {
 port number;
 }
}
server-inactivity-timeout minutes;
}
```

The preceding configuration settings that are described define a TWAMP server on the router that enables a TWAMP client to connect to the server using any media interface IP address such as a **ge-** interface. In such a scenario, the router functions as a TWAMP server and timestamping is performed in the ukernel of the media-facing FPC.

To configure an inline TWAMP server, which causes timestamping to be performed as part of the inline services (**si-**) interface processing, configure the amount of bandwidth reserved on each Packet Forwarding Engine for tunnel traffic using inline services by including the **bandwidth (1g | 10g)** statement at the **[edit chassis fpc slot-number pic number inline-services]** hierarchy level and specify the service PIC logical interface that provides the TWAMP service by including the **twamp-server** statement at the **[edit interfaces sp-fpc/pic/port unit logical-unit- number family inet]** hierarchy level.

- To specify the list of allowed control client hosts that can connect to this server, include the **client-list** statement at the **[edit services rpm twamp server]** hierarchy level. Each value you include must be a Classless Interdomain Routing (CIDR) address (IP address plus mask) that represents a network of allowed hosts. You can include multiple client lists, each of which can contain a maximum of 64 entries. You must configure at least one client address to enable TWAMP.
- You must specify the authentication mode by including the **authentication-mode** statement at the **[edit services rpm twamp server]** hierarchy level. There is no default value. You can configure **authenticated** or **encrypted** mode, based on RFC 4656; if there is no authentication or encryptions mode specified, you should set the value to **none**. This statement is required in the TWAMP configuration.
- To specify the inactivity timeout period in seconds, include the **inactivity-timeout** statement at the **[edit services rpm twamp server]** hierarchy level. By default, the value is **1800**; the range is 0 through 3600 seconds.
- To specify the maximum number of concurrent connections the server can have to client hosts, include the **maximum-connections** statement at the **[edit services rpm twamp server]** hierarchy level. The allowed range of values is 1 through 1000 and the default value is 64. You can also limit the number of connections the server can make to a particular client host by including the **maximum-connections-per-client** statement. The allowed range of values is 1 through 500 and the default value is 64.
- To specify the maximum number of sessions the server can have running at one time, include the **maximum-sessions** statement at the **[edit services rpm twamp server]** hierarchy level. The allowed range of values is 1 through 2048 and the default value is

64. You can also limit the number of sessions the server can have on a single connection by including the **maximum-sessions-per-connection** statement.

- To specify the TWAMP server listening port, include the **port** statement at the **[edit services rpm twamp server]** hierarchy level. The range is 1 through 65,535.
- To specify the server inactivity timeout period in minutes, include the **server-inactivity-timeout** statement at the **[edit services rpm twamp server]** hierarchy level. The range is 0 through 30 minutes.
- To specify the TWAMP servers on specific routing instances, instead of associating the TWAMP server at the system-level to apply to all routing instances configured on a router, include the **routing-instance-list instance-name port port-number** statement at the **[edit services rpm twamp server]** hierarchy level. The port number of the specified routing instance is used for TWAMP probes that are received by a TWAMP server. The default routing instance is Internet routing table inet.0. If you do not specify a routing instance, the TWAMP probe applies to all routing instances. To apply the TWAMP probe to only the default routing instance, you must explicitly set the value of instance-name to default. If an interface is not part of any routing instance, the default port is used for TWAMP probes. You can configure up to 100 routing instances for a TWAMP server.

## Configuring BGP Neighbor Discovery Through RPM

BGP neighbors can be configured at the following hierarchy levels:

- **[edit protocols bgp group group-name]**—Default logical system and default routing instance.
- **[edit routing-instances instance-name protocols bgp group group-name]**—Default logical system with a specified routing instance.
- **[edit logical-systems logical-system-name protocols bgp group group-name]**—Configured logical system and default routing instance.
- **[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group group-name]**—Configured logical system with a specified routing instance.

When you configure BGP neighbor discovery through RPM, if you do not specify a logical system, the RPM probe applies to configured BGP neighbors for all logical systems. If you do not specify a routing instance, the RPM probe applies to configured BGP neighbors in all routing instances. You can explicitly configure RPM probes to apply only to the default logical system, the default routing instance, or to a particular logical system or routing instance.

To configure BGP neighbor discovery through RPM, configure the probe properties at the **[edit services rpm bgp]** hierarchy:

```
data-fill data;
data-size size;
destination-port port;
history-size size;
logical-system logical-system-name [routing-instances routing-instance-name];
```

**moving-average-size** *number*;  
**probe-count** *count*;  
**probe-interval** *seconds*;  
**probe-type** *type*;  
**routing-instances** *instance-name*;  
**test-interval** *interval*;

- To specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes, include the **data-fill** statement at the **[edit services rpm bgp]** hierarchy level. The value can be a hexadecimal value.
- To specify the size of the data portion of ICMP probes, include the **data-size** statement at the **[edit services rpm bgp]** hierarchy level. The size can be from **0** through **65400** and the default size is **0**.
- To specify the User Datagram Protocol (UDP) port or Transmission Control Protocol (TCP) port to which the probe is sent, include the **destination-port** statement at the **[edit services rpm bgp]** hierarchy level. The **destination-port** statement is used only for the UDP and TCP probe types. The value can be **7** or from **49160** through **65535**.
- To specify the number of stored history entries, include the **history-size** statement at the **[edit services rpm bgp]** hierarchy level. Specify a value from **0** to **512**. The default is **50**.
- To specify the logical system used by ICMP probes, include the **logical-system** *logical-system-name* statement at the **[edit services rpm bgp]** hierarchy level. If you do not specify a logical system, the RPM probe applies to configured BGP neighbors for all logical systems. To apply the probe to only the default logical system, you must set the value of *logical-system-name* to null.
- To specify a number of samples for making statistical calculations, include the **moving-average-size** statement at the **[edit services rpm bgp]** hierarchy level. Specify a value from 0 through 255.
- To specify the number of probes within a test, include the **probe-count** statement at the **[edit services rpm bgp]** hierarchy level. Specify a value from 1 through 15.
- To specify the time to wait between sending packets, include the **probe-interval** statement at the **[edit services rpm bgp]** hierarchy level. Specify a value from 1 through 255 seconds.
- To specify the packet and protocol contents of the probe, include the **probe-type** statement at the **[edit services rpm bgp]** hierarchy level. The following probe types are supported:
  - **icmp-ping**—Sends ICMP echo requests to a target address.
  - **icmp-ping-timestamp**—Sends ICMP timestamp requests to a target address.
  - **tcp-ping**—Sends TCP packets to a target.
  - **udp-ping**—Sends UDP packets to a target.
  - **udp-ping-timestamp**—Sends UDP timestamp requests to a target address.





**NOTE:** Some probe types require additional parameters to be configured. For example, when you specify the `tcp-ping` or `udp-ping` option, you must configure the destination port using the `destination-port port` statement. The `udp-ping-timestamp` option requires a minimum data size of 12; any smaller data size results in a commit error. The minimum data size for TCP probe packets is 1.

- To specify the routing instance used by ICMP probes, include the **routing-instances** statement at the `[edit services rpm bgp]` hierarchy level. The default routing instance is Internet routing table `inet.0`. If you do not specify a routing instance, the RPM probe applies to configured BGP neighbors in all routing instances. To apply the RPM probe to only the default routing instance, you must explicitly set the value of *instance-name* to **default**.
- To specify the time to wait between tests, include the **test-interval** statement at the `[edit services bgp probe]` hierarchy level. Specify a value from 1 through 86400 seconds.



**NOTE:** Starting with Junos OS Release 15.1, the minimum period for which the RPM client waits between two tests is modified to be 1 second instead of 0 seconds. Also, if you do not configure the test interval, the default value is 1 second.

#### Related Documentation

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)
- [Examples: Configuring BGP Neighbor Discovery Through RPM on page 973](#)

## Examples: Configuring BGP Neighbor Discovery Through RPM

Configure BGP neighbor discovery through RPM for all logical systems and all routing instances:

```
[edit services rpm]
bgp {
 probe-type icmp-ping;
 probe-count 5;
 probe-interval 1;
 test-interval 60;
 history-size 10;
 data-size 255;
 data-fill 0123456789;
}
```

Configure BGP neighbor discovery through RPM for only the following logical systems and routing instances: **LS1/RI1**, **LS1/RI2**, **LS2**, and **RI3**:

```
[edit services rpm]
bgp {
```

```
probe-type icmp-ping;
probe-count 5;
probe-interval 1;
test-interval 60;
history-size 10;
data-size 255;
data-fill 0123456789;
logical-system {
 LS1 {
 routing-instances {
 RI1;
 RI2;
 }
 }
 LS2;
}
routing-instance {
 RI3;
}
}
```



**NOTE:** The `logical-system` statement is not supported on PTX Series Packet Transport Routers.

Configure BGP neighbor discovery through RPM for only the default logical system and default routing instance:

```
[edit services rpm]
bgp {
 probe-type icmp-ping;
 probe-count 5;
 probe-interval 1;
 test-interval 60;
 history-size 10;
 data-size 255;
 data-fill 0123456789;
 logical-system {
 null {
 routing-instances {
 default;
 }
 }
 }
}
```

**Related  
Documentation**

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [Configuring BGP Neighbor Discovery Through RPM on page 971](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)

## Tracing RPM Operations

Tracing operations track all RPM operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

By default, no events are traced. If you include the **traceoptions** statement at the **[edit services rpm]** hierarchy level, the default tracing behavior is the following:

- Important events are logged in a file called **rmopd** located in the **/var/log** directory.
- When the log file reaches 128 kilobytes (KB), it is renamed **rmopd.0**, then **rmopd.1**, and so on, until there are three trace files. Then the oldest trace file (**rmopd.2**) is overwritten. (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)
- Log files can be accessed only by the user who configures the tracing operation.

You can change this default behavior by using the **traceoptions** statements. Changing the defaults is described in the following sections:

1. [Configuring the RPM Log File Name on page 1153](#)
2. [Configuring the Number and Size of RPM Log Files on page 1153](#)
3. [Configuring Access to the Log File on page 1154](#)
4. [Configuring a Regular Expression for Lines to Be Logged on page 1154](#)
5. [Configuring the Trace Operations on page 1154](#)

### Configuring the RPM Log File Name

By default, the name of the file that records RPM trace output is **rmopd**. To specify a different file name:

```
[edit services rpm traceoptions]
user @host set file filename
```

### Configuring the Number and Size of RPM Log Files

To configure the limits on the number and size of RPM trace files:

```
[edit services rpm traceoptions]
user@host set file filename files number size size
```

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

For example, set the maximum file size to 2 MB, and the maximum number of files to 20 for a log file named **rpmtrace**:

```
[edit services rpm traceoptions]
user@host set file rpmtrace files 20 size 2MB
```

When the **rpmtrace** file reaches 2 MB, it is renamed **rpmtrace.0**, and a new file called **rpmtrace** is created. When the new **rpmtrace** reaches 2 MB, **rpmtrace.0** is renamed

**rpmtrace.1** and **rpmtrace** is renamed **rpmtrace.0**. This process repeats until there are 20 trace files. Then the oldest file (**rpmtrace.19**) is overwritten by **rpmtrace.18**.

## Configuring Access to the Log File

By default, log files can be accessed only by the user who configures the tracing operation.

To specify that any user can read all log files:

```
[edit services rpm traceoptions]
user@host set file filename world-readable
```

To explicitly set the default behavior:

```
[edit services rpm traceoptions]
user@host set file filename no-world-readable
```

## Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

To refine the output by specifying a regular expression (regex) to be matched:

```
[edit services rpm traceoptions]
user@host set file filename match regular-expression
```

## Configuring the Trace Operations

By default, if the **traceoptions** configuration is present, only important events are logged. You can configure the trace operations to be logged by including the following statements at the **[edit services rpm traceoptions]** hierarchy level:

```
flag {
 all;
 configuration;
 error;
 ipc;
 ppm;
 statistics
}
```

[Table 7 on page 51](#) describes the meaning of the RPM tracing flags.

**Table 46: RPM Tracing Flags**

Flag	Description	Default Setting
<b>all</b>	Trace all operations.	Off
<b>configuration</b>	Trace configuration events.	Off
<b>error</b>	Trace events related to catastrophic errors in daemon.	Off
<b>ipc</b>	Trace IPC events.	Off

Table 46: RPM Tracing Flags (*continued*)

Flag	Description	Default Setting
<b>ppm</b>	Trace ppm events.	Off
<b>statistics</b>	Trace statistics.	Off

## Examples: Configuring Real-Time Performance Monitoring

Configure an RPM instance identified by the probe name **probe1** and the test name **test1**:

```
[edit services rpm]
probe probe1 {
 test test1 {
 dscp-code-points 001111;
 probe-interval 1;
 probe-type icmp-ping;
 target address 172.17.20.182;
 test-interval 20;
 thresholds rtt 10;
 traps rtt-exceeded;
 }
}
probe-server {
 tcp {
 destination-interface lt-0/0/0.0
 port 50000;
 }
 udp {
 destination-interface lt-0/0/0.0
 port 50001;
 }
}
probe-limit 200;
```

Configure packet classification, using **lt-** interfaces to send the probe packets to a logical tunnel input interface. By sending the packet to the logical tunnel interface, you can configure regular and multifield classifiers, firewall filters, and header rewriting for the probe packets. To use the existing tunnel framework, the **dlci** and **encapsulation** statements must be configured.

```
[edit services rpm]
probe p1 {
 test t1 {
 probe-type icmp-ping;
 target address 10.8.4.1;
 probe-count 10;
 probe-interval 10;
 test-interval 10;
 source-address 10.8.4.2;
 dscp-code-points ef;
 data-size 100;
 destination-interface lt-0/0/0.0;
 }
}
```

```
}
[edit interfaces]
lt-0/0/0 {
 unit 0 {
 encapsulation frame-relay;
 dlci 10;
 peer-unit 1;
 family inet;
 }
 unit 1 {
 encapsulation frame-relay;
 dlci 10;
 peer-unit 0;
 family inet;
 }
}
[edit class-of-service]
interfaces {
 lt-0/0/0 {
 unit 1 {
 classifiers {
 dscp default;
 }
 }
 }
}
}
```

Configure an input filter on the interface on which the RPM probes are received. This filter enables prioritization of the received RPM packets, separating them from the regular data packets received on the same interface.

```
[edit firewall]
filter recos {
 term recos {
 from {
 source-address {
 10.8.4.1/32;
 }
 destination-address {
 10.8.4.2/32;
 }
 }
 then {
 loss-priority high;
 forwarding-class network-control;
 }
 }
}
[edit interfaces]
fe-5/0/0 {
 unit 0 {
 family inet {
 filter {
 input recos;
 }
 address 10.8.4.2/24;
 }
 }
}
```

```

 }
 }
}

```

Configure an RPM instance and enable RPM for the extension-provider packages on the adaptive services interface:

```

[edit services rpm]
probe probe1 {
 test test1 {
 data-size 1024;
 data-fill 0;
 destination-interface ms-1/2/0.10;
 dscp-code-points 001111;
 probe-count 10;
 probe-interval 1;
 probe-type icmp-ping;
 target address 172.17.20.182;
 test-interval 20;
 thresholds rtt 10;
 traps rtt-exceeded;
 }
}
[edit interfaces]
ms-1/2/0 {
 unit 0 {
 family inet;
 }
 unit 10 {
 rpm client;
 family inet {
 address 1.1.1.1/32;
 }
 }
}
[edit chassis]
fpc 1 {
 pic 2 {
 adaptive-services {
 service-package {
 extension-provider {
 control-cores 1;
 data-cores 1;
 object-cache-size 512;
 policy-db-size 64;
 package jservices-rpm;
 syslog {
 daemon any;
 }
 }
 }
 }
 }
}
}
}
}

```



**NOTE:** TWAMP is not supported on PTX Series Packet Transport Routers.

Configure the minimum statements necessary to enable TWAMP:

```
[edit services]
rpm {
 twamp {
 server {
 authentication-mode none;
 port 10000; # Twamp server's listening port
 client-list LIST-1 { # LIST-1 is the name of the client-list. Multiple lists can be
 configured.
 address {
 20.0.0.2/30; # IP address of the control client.
 }
 }
 }
 }
}
[edit interfaces sp-5/0/0]
unit 0 {
 family inet;
}
unit 10 {
 rpm {
 twamp-server; # You must configure a separate logical interface on the service PIC
 interface for the TWAMP server.
 }
 family inet {
 address 50.50.50.50/32; # This address must be a host address with a 32-bit mask.
 }
}
[edit chassis]
fpc 5 {
 pic 0 {
 adaptive-services {
 service-package layer-2; # Configure the service PIC to run in Layer 2 mode.
 }
 }
}
```

Configure additional TWAMP settings:

```
[edit services]
rpm {
 twamp {
 server {
 maximum-sessions 5;
 maximum-sessions-per-connection 2;
 maximum-connections 3;
 maximum-connections-per-client 1;
 port 10000;
 server-inactivity-timeout ;
 client-list LIST-1 {
 address {
```



```

 20.0.0.2/30;
 }
}
}
}
}

```

#### Related Documentation

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)
- [Examples: Configuring BGP Neighbor Discovery Through RPM on page 973](#)

## Enabling RPM for the Junos OS extension-provider package

Real-time performance monitoring (RPM), which has been supported on the adaptive services interface, is now supported by the Junos OS extension-provider package. RPM is supported on all platforms and service PICs that support the extension-provider package.



**NOTE:** In Junos OS releases earlier than 12.3, the extension provider package was variously known as MP-SDK, Junos Services Framework (JSF), and eJunos.

To enable RPM for the Junos OS extension-provider package on the adaptive services interface, configure the **object-cache-size**, **policy-db-size**, and **package** statements at the **[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]** hierarchy level. For the extension-provider package, **package-name** in the **package package-name** statement is **jservices-rpm**.

For more information about the extension-provider package, see the *SDK Applications Configuration Guide and Command Reference*.

The following example shows how to enable RPM for the extension-provider package on the adaptive services interface:

```

chassis fpc 1 {
 pic 2 {
 adaptive-services {
 service-package {
 extension-provider {
 control-cores 1;
 data-cores 1;
 object-cache-size 512;
 policy-db-size 64;
 package jservices-rpm;
 syslog daemon any;
 }
 }
 }
 }
}

```

}

**Related  
Documentation**

- [Real-Time Performance Monitoring Services Overview on page 957](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)
- [Examples: Configuring Real-Time Performance Monitoring on page 977](#)
- [destination-interface on page 1652](#)

# Testing the Performance of Network Devices Using RFC 2544-Based Benchmarking

- [RFC2544-Based Benchmarking Tests Overview on page 1161](#)
- [Configuring an RFC 2544-Based Benchmarking Test on page 1163](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for Layer 3 IPv4 Services on page 1168](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for UNI Direction of Ethernet Pseudowires on page 1175](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for NNI Direction of Ethernet Pseudowires on page 1183](#)

## RFC2544-Based Benchmarking Tests Overview

RFC2544 defines a series of tests that can be used to describe the performance characteristics of a network-interconnecting device, such as a router, and outlines specific formats to report the results of the tests. These tests can be used to benchmark interconnected network devices and devise a guideline or a measurement pattern to analyze the health and efficiency of the network devices. These tests are the standard benchmarking tests for Ethernet networks and are known as RFC2544-based benchmarking tests. These tests measure throughput, latency, frame loss rate, and bursty frames. The test methodology enables you to define various parameters such as different frame sizes to be examined (64, 128, 256, 512, 1024, 1280, and 1518 bytes), the test time for each test iteration (10 seconds to 1,728,000 seconds), and the frame format (UDP-over-IP).



**NOTE:** RFC2544-based benchmarking tests support only UDP over IPv4 test traffic.

An RFC2544-based benchmarking test is performed by transmitting test packets from a device that functions as the generator or the initiator (which is also called the originator). These packets are sent to a device that functions as a reflector, which receives and returns the packets to the initiator.

Juniper Networks MX104 3D Universal Edge Routers support only the reflector function and the corresponding benchmarking tests. These tests display only the reflecting benchmarking tests. These benchmarking tests display the results of the test. For instance, in the case of the throughput test, the results display the number of transmitted frames and the number of received frames.

The RFC2544-based benchmarking test methodology assesses different parameters that are defined in service-level agreements (SLAs). By measuring the performance availability, transmission delay, link bursts, and service integrity, a carrier provider can certify that the working parameters of the deployed Ethernet circuit comply with the SLA and other defined policies.

[Table 37 on page 984](#) describes the different network topologies in which the benchmarking test is supported.

**Table 47: Supported Network Topologies for RFC2544 Benchmarking Tests**

Service Type	Traffic Direction	Mode	Initial Release on MX104 Routers	Whether the Benchmarking Test Is Supported
E-Line and E-LAN (family <b>bridge</b> )	(UNI) Egress	Port Port, VLAN	14.2R1 (E-Line and E-LAN family bridge)	Supported
E-Line (family <b>ccc</b> )	Ingress Egress		13.3R1 (E-Line Pseudowire)	Supported
IP Services (family <b>inet</b> )	NNI		13.3R1	Supported



**NOTE:** You can configure a total of four simultaneous active reflection sessions. The four active reflection sessions can be of the same type or can be a combination of the different types of reflection sessions. For instance, you can configure either four IPv4 reflection sessions or two pseudowire reflection sessions, one Layer 2 reflection session, and one IPv4 reflection session. The maximum reflection bandwidth supported is 4Gbps.

[Table 38 on page 985](#) lists the interfaces and the reflection type on which the benchmarking tests are supported.

**Table 48: Supported Interfaces for RFC2544 Benchmarking Tests**

Type of Reflection	Gigabit Interfaces (ge)	Aggregated Interfaces (ae)	10G Interfaces (xe)	Pseudo Interfaces (irb, lt, vt, lo0, and others)
IPv4	Yes	No	No	No
Pseudowire Ingress	Yes	No	No	No
Pseudowire Egress	Yes	No	No	No
Layer 2 Bridge	Yes	Yes	Yes	No

All active RFC2544-based benchmarking tests are stopped when any of the following events takes place either in the initiator or in the reflector:

- System events such as Packet Forwarding Engine restarts, routing engine restarts, and so on.
- Test interface change events such as deactivation and reactivation of the interface, disabling and enabling of the interface, and so on.

After the benchmarking tests are stopped, the test states of the tests are removed and the user can restart the same test. Other ongoing tests on other interfaces are not interrupted.



**NOTE:** RFC2544-based benchmarking tests are not supported during unified in-service software upgrade (ISSU) and graceful routing engine switchover (GRES).

#### Related Documentation

- [Configuring an RFC 2544-Based Benchmarking Test on page 989](#)
- [Supported RFC2544-Based Benchmarking Statements on MX104 Routers on page 988](#)

## Configuring an RFC 2544-Based Benchmarking Test

You can configure a benchmarking test to detect and measure performance attributes, such as throughput, latency, frame loss, and bursty or back-to-back frames, of network devices. RFC 2544-based benchmarking test is performed by transmitting test packets from a device that functions as the generator or the initiator. These packets are sent to a device that functions as the reflector, which receives and returns the packets back to the initiator.



**NOTE:** The test configuration is applied only when you start the test. If you update the test configuration during the test, you have to start the test again for the updated configuration to take effect.

You must configure a test profile and reference the test profile in a unique test name that defines the parameters for the test to be performed on a certain device. However, the test profile is required when the test mode is configured as initiation and termination. The **test-profile** parameter is disregarded when the test mode is configured as reflection. MX104 routers support only the reflection function in the RFC 2544-based benchmarking tests. A reflection service does not use the parameters specified in the test profile.

The following topics describe how to configure a test name for an RFC 2544-based benchmarking test on an MX104 router for Layer 3 IPv4, Ethernet pseudowire, and Layer 2 bridge networks:

- [Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a IPv4 Network on page 1164](#)
- [Configuring a Test Name for an RFC 2544-Based Benchmarking Test for an Ethernet Pseudowire on page 1165](#)
- [Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a Layer 2 E-LAN Service in Bridge Domain on page 1167](#)

## Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a IPv4 Network

You can configure a test name by including the **test-name test-name** statement at the **[edit services rpm rfc2544-benchmarking]** hierarchy level. In the test name, you can configure attributes of the test iteration, such as the address family (type of service, IPv4 or Ethernet), the logical interface, and test duration that are used for a benchmarking test to be run.

To configure a test name and define its attributes for an IPv4 network:

1. In configuration mode, go to the **[edit services]** hierarchy level.

```
[edit]
user@host# edit services
```

2. Configure a instance.

```
[edit services]
user@host# edit rpm
```

3. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

4. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

5. Specify the test mode for the packets that are sent during the benchmarking test. The **reflect** option causes the test frames to be reflected on the IPv4 network.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```

6. Configure the address type family for the benchmarking test. The **inet** option indicates that the test is run on an IPv4 service.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family inet
```

7. Configure the destination IPv4 address for the test packets. This parameter is required only if you configure IPv4 family **inet**. If you do not configure the destination IPv4 address, the default value of 192.168.1.20 is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set destination-ipv4-address address
```

8. Specify the UDP port of the destination to be used in the UDP header for the generated frames. If you do not specify the UDP port, the default value of 4041 is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set destination-udp-port port-number
```

9. (Optional) Specify the source IPv4 address to be used in generated test frames. If you do not configure the source IPv4 address for **inet** family, the source address of the interface is used to transmit the test frames.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set source-ipv4-address address
```

10. Specify the UDP port of the source to be used in the UDP header for the generated frames. If you do not specify the UDP port, the default value of 4041 is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set source-udp-port port-number
```

11. Specify the logical interface on which the RFC 2544-based benchmarking test is run. If you configure an **inet** family and the test mode to reflect the frames back on the sender from the other end, then the logical interface is used as the interface to enable the reflection service (reflection is performed on the packets entering the specified interface). If you not configure the logical interface for reflection test mode, then a lookup is performed on the source IPv4 address to determine the interface that hosts the address.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface interface-name
```

## Configuring a Test Name for an RFC 2544-Based Benchmarking Test for an Ethernet Pseudowire

You can configure a test name by including the **test-name test-name** statement at the **[edit services rpm rfc2544-benchmarking]** hierarchy level. In the test name, you can configure attributes of the test iteration, such as the address family (type of service IPv4 or Ethernet), the logical interface, and test duration, that are used for a benchmarking test to be run. The test name combined with the test profile represent a single real-time performance monitoring (RPM) configuration instance.

To configure a test name and define its attributes for an Ethernet Pseudowire:

1. In configuration mode, go to the **[edit services]** hierarchy level.

```
[edit]
user@host# edit services
```

2. Configure an RPM service instance.

```
[edit services]
user@host# edit rpm
```

3. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

4. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

5. Specify the test mode for the packets that are sent during the benchmarking test. The **reflect** option causes the test frames to be reflected on the Ethernet pseudowire.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```

6. Configure the address type family for the benchmarking test. The **ccc** option indicates that the test is run on a CCC or Ethernet pseudowire service.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```

7. Specify the direction of the interface on which the test must be run. This parameter is valid only for a family. To enable the test to be run in the egress direction of the interface (network-to-network interface (NNI)), use the **egress** option. To enable the test to be run in the ingress direction of the interface (user-to-network interface (UNI)), use the **ingress** option.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction egress
```

8. (Optional) Specify the source IPv4 address to be used in generated test frames. If you do not configure the source IPv4 address for family, the default value of 192.168.1.10 is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set source-ipv4-address address
```

9. Specify the logical interface on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface interface-name
```



## Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a Layer 2 E-LAN Service in Bridge Domain

You can configure a test name by including the **test-name test-name** statement at the **[edit services rpm rfc2544-benchmarking]** hierarchy level. In the test name, you can configure attributes of the test iteration, such as the address family (bridge), the logical interface, and test duration, that are used for a benchmarking test to be run. The test name combined with the test profile represent a single real-time performance monitoring (RPM) configuration instance.

To configure a test name and define its attributes for a layer 2 E-LAN service in Bridge domains:

1. In configuration mode, go to the **[edit services]** hierarchy level.

```
[edit]
user@host# edit services
```

2. Configure an RPM service instance.

```
[edit services]
user@host# edit rpm
```

3. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

4. Define a name for the test—for example, l2b-test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name l2b-test1
```

5. Specify the source and destination MAC addresses of the test packet. Both these parameters are valid only for the bridge family.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set source-mac-address address destination-mac-address address
```

6. Specify the service type under test. This parameter is applicable only for the bridge family.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set service-type elan
```

7. Specify the test mode for the packets that are sent during the benchmarking test. The **reflect** option causes the test frames to be reflected over the Layer 2 bridge.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set mode reflect
```

8. Configure the address type family for the benchmarking test. The **bridge** option indicates that the test is run on a E-LAN service over a bridge domain.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set family bridge
```

9. Specify the direction of the interface on which the test must be run. This parameter is valid only for a family. To enable the test to be run in the egress direction of the interface (network-to-network interface (NNI)), use the **egress** option. To enable the test to be run in the ingress direction of the interface (user-to-network interface (UNI)), use the **ingress** option.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set direction egress
```

10. Specify the logical interface on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set test-interface interface-name
```

#### Related Documentation

- [RFC2544-Based Benchmarking Tests Overview on page 983](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for UNI Direction of Ethernet Pseudowires on page 1001](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for NNI Direction of Ethernet Pseudowires on page 1008](#)
- [Example: Configuring RFC2544-Based Benchmarking Tests for Layer 2 E-LAN Services in Bridge Domains on page 1016](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for Layer 3 IPv4 Services on page 993](#)

### Example: Configuring an RFC 2544-Based Benchmarking Test for Layer 3 IPv4 Services

- [Requirements on page 1168](#)
- [Overview on page 1168](#)
- [Configuration on page 1169](#)
- [Verifying the Results of the Benchmarking Test for Layer 3 IPv4 Services on page 1175](#)

#### Requirements

This example uses the following hardware and software components:

- 
- An ACX Series Universal Access Router—
- Junos OS Release or later

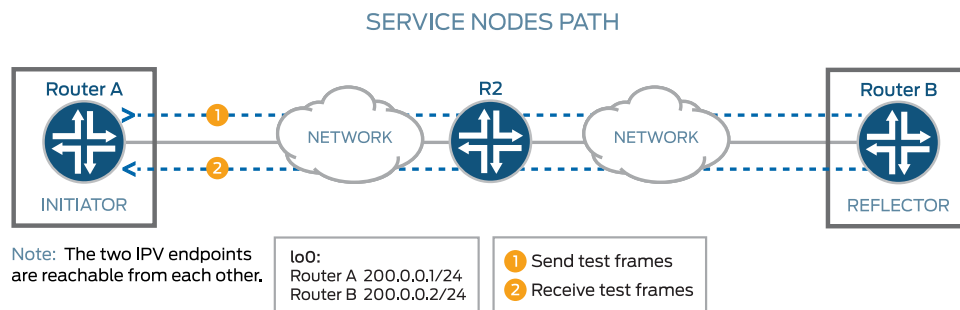
#### Overview

Consider a sample topology in which a router, Router A, functions as an initiator and terminator of the test frames for an RFC 2544-based benchmarking test. Router A is connected over a Layer 3 network to another router, Router B, which functions as a reflector to reflect back the test frames it receives from Router A. IPv4 is used for transmission of test frames over the Layer 3 network. This benchmarking test is used to compute the IPv4 service parameters between Router A and Router B. Logical interfaces

on both the routers are configured with IPv4 addresses to measure the performance attributes, such as throughput, latency, frame loss, and bursty frames, of network devices for the IPv4 service.

Figure 37 on page 994 shows the sample topology to perform an RFC 2544 test for a Layer 3 IPv4 service.

Figure 43: RFC 2544-Based Benchmarking Test for a Layer 3 IPv4 Service



## Configuration

In this example, you configure the benchmarking test for a Layer 3 IPv4 service that is between interface ge-0/0/0 on Router A and interface ge-0/0/4 on Router B to detect and analyze the performance of the interconnecting routers.

- [Configuring Benchmarking Test Parameters on Router A on page 1170](#)
- [Configuring Benchmarking Test Parameters on Router B on page 1172](#)
- [Results on page 1174](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

### Configuring Benchmarking Test Parameters on Router A

```
set interfaces ge-0/0/0 unit 0 family inet address 200.0.0.1/24
set interfaces ge-0/0/0 unit 0 family mpls
set rfc2544-benchmarking profiles test-profile throughput test-type throughput
set rfc2544-benchmarking profiles test-profile throughput packet-size 64
set rfc2544-benchmarking profiles test-profile throughput test-duration 20m
set rfc2544-benchmarking profiles test-profile throughput bandwidth-kbps 500
set rfc2544-benchmarking tests test-name test1 test-profile throughput
set rfc2544-benchmarking tests test-name test1 interface ge-0/0/0.1
set rfc2544-benchmarking tests test-name test1 mode initiate-and-terminate
set rfc2544-benchmarking tests test-name test1 family inet
set rfc2544-benchmarking tests test-name test1 dest-address 200.0.0.2
set rfc2544-benchmarking tests test-name test1 udp-port 4001
```

## Configuring Benchmarking Test Parameters on Router B

```
set interfaces ge-0/0/4 unit 0 family inet address 200.0.0.2/24
set interfaces ge-0/0/4 unit 0 family mpls
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/4.1
set services rpm rfc2544-benchmarking tests test-name test1 mode reflect
set services rpm rfc2544-benchmarking tests test-name test1 family inet
set services rpm rfc2544-benchmarking tests test-name test1 dest-address 200.0.0.1
set services rpm rfc2544-benchmarking tests test-name test1 udp-port 4001
```

---

### Configuring Benchmarking Test Parameters on Router A

#### Step-by-Step Procedure

The following require you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the test parameters on Router A:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:  

```
[edit]
user@host# edit interfaces
```
2. Configure the interface on which the test must be run.  

```
[edit interfaces]
user@host# edit ge-0/0/0
```
3. Configure a logical unit and specify the protocol family.  

```
[edit interfaces ge-0/0/0]
user@host# edit unit 0 family inet
```
4. Specify the address for the logical interface.  

```
[edit interfaces ge-0/0/0 unit 0 family inet]
user@host# set address 200.0.0.1/24
```
5. Enter the **up** command to go the previous level in the configuration hierarchy.  

```
[edit interfaces ge-0/0/0 unit 0 family inet]
user@host# up
```
6. Configure the MPLS family on the logical interface.  

```
[edit interfaces ge-0/0/0 unit 0]
user@host# set family mpls
```
7. Go to the top level of the configuration command mode.  

```
[edit interfaces ge-0/0/0 unit 0]
user@host# top
```
8. In configuration mode, go to the **[edit services]** hierarchy level.  

```
[edit]
user@host# edit services
```

9. Configure a real-time performance monitoring service (RPM) instance.  
[edit services]  
user@host# **edit rpm**
10. Configure an RFC 2544-based benchmarking test for the RPM instance.  
[edit services rpm]  
user@host# **edit rfc2544-benchmarking**
11. Define a name for a test profile—for example, throughput.  
[edit services rpm rfc2544-benchmarking]  
user@host# **edit profiles test-profile throughput**
12. Configure the type of test to be performed as throughput.  
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# **set test-type throughput**
13. Specify the size of the test packet as 64 bytes.  
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# **set test-type packet-size 64**
14. Specify the period for which the test is to be performed in hours, minutes, or seconds by specifying a number followed by the letter h (for hours), m (for minutes), or s (for seconds), respectively.  
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# **set test-type test-duration 20m**
15. Define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.  
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# **set test-type bandwidth-kbps 500**
16. Enter the **up** command to go the previous level in the configuration hierarchy.  
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# **up**
17. Enter the **up** command to go the previous level in the configuration hierarchy.  
[edit services rpm rfc2544-benchmarking profiles]  
user@host# **up**
18. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.  
[edit services rpm rfc2544-benchmarking]  
user@host# **edit tests test-name test1**
19. Specify the name of the test profile—for example, throughput—to be associated with a particular test name.  
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# **set test-profile throughput**
20. Specify the logical interface, ge-0/0/0.1, on which the RFC 2544-based benchmarking test is run.  
[edit services rpm rfc2544-benchmarking tests test-name test1]

```
user@host# set test-interface ge-0/0/0.1
```

21. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode initiate-and-terminate
```

22. Configure the address type family, **inet**, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family inet
```

23. Configure the destination IPv4 address for the test packets.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set dest-address 200.0.0.2
```

24. Specify the UDP port of the destination to be used in the UDP header for the generated frames as 4001.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set udp-port 4001
```

25. Start the benchmarking test on the initiator.

```
user@> test services rpm rfc2544-benchmarking test test1 start
```

After the test is successfully completed, it is automatically stopped at the initiator.

---

### Configuring Benchmarking Test Parameters on Router B

#### Step-by-Step Procedure

The following you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the test parameters on Router B:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

```
[edit]
user@host# edit interfaces
```

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@host# edit ge-0/0/4
```

3. Configure a logical unit and specify the protocol family as **inet**.

```
[edit interfaces ge-0/0/4]
user@host# edit unit 0 family inet
```

4. Specify the address for the logical interface.

```
[edit interfaces ge-0/0/4 unit 0 family inet]
user@host# set address 200.0.0.2/24
```

5. Enter the **up** command to go the previous level in the configuration hierarchy.

```
[edit interfaces ge-0/0/4 unit 0 family inet]
user@host# up
```

6. Configure the MPLS family on the logical interface.  

```
[edit interfaces ge-0/0/4 unit 0]
user@host# set family mpls
```
7. Go to the top level of the configuration command mode.  

```
[edit interfaces ge-0/0/4 unit 0]
user@host# top
```
8. In configuration mode, go to the **[edit services]** hierarchy level.  

```
[edit]
user@host# edit services
```
9. Configure a real-time performance monitoring service (RPM) instance.  

```
[edit services]
user@host# edit rpm
```
10. Configure an RFC 2544-based benchmarking test for the RPM instance.  

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```
11. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.  

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```
12. Specify the logical interface, ge-0/0/4.1, on which the RFC 2544-based benchmarking test is run.  

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface ge-0/0/4.1
```
13. Specify **reflect** as the test mode for the packets that are sent during the benchmarking test.  

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```
14. Configure the address type family, **inet**, for the benchmarking test.  

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family inet
```
15. Configure the destination IPv4 address for the test packets as 200.0.0.1.  

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set dest-address 200.0.0.1
```
16. Specify the UDP port of the destination to be used in the UDP header for the generated frames as 4001.  

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set udp-port 4001
```
17. Start the benchmarking test on the reflector.  

```
user@host> test services rpm rfc2544-benchmarking test test1 start
```

After the test is successfully completed at the initiator, you can stop the test at the reflector by entering the **test services rpm rfc2544-benchmarking test test1** command.

## Results

---

In configuration mode, confirm your configuration on Router A and Router B by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

### Benchmarking Test Parameters on Router A:

```
[edit interfaces]
ge-0/0/0 {
 unit 0 {
 family inet {
 address 200.0.0.1/24;
 }
 family mpls;
 }
}

[edit services rpm]
rfc2544-benchmarking {
 profiles {
 test-profile throughput {
 test-type throughput
 packet-size 64;
 test-duration 20m;
 bandwidth-kbps 500;
 }
 }

 tests {
 test-name test1 {
 test-profile throughput;
 interface ge-0/0/0.1;
 mode initiate,terminate;
 family inet;
 dest-address 200.0.0.2
 udp-port 4001;
 }
 }
}
```

### Benchmarking Test Parameters on Router B:

```
[edit interfaces]
ge-0/0/4 {
 unit 0 {
 family inet {
 address 200.0.0.2/24;
 }
 family mpls;
 }
}

[edit services rpm]
rfc2544-benchmarking {
 # Note, When in reflector mode, test profile is not needed
```



```
tests {
 test-name test1 {
 interface ge-0/0/4.1;
 mode reflect;
 family inet;
 dest-address 200.0.0.1;
 udp-port 4001;
 }
}
```

After you have configured the device, enter the **commit** command in configuration mode.

## Verifying the Results of the Benchmarking Test for Layer 3 IPv4 Services

Examine the results of the benchmarking test that is performed on the configured service between Router A and Router B.

- [Verifying the Benchmarking Test Results on page 1175](#)

---

### Verifying the Benchmarking Test Results

<b>Purpose</b>	Verify that the necessary and desired statistical values are displayed for the benchmarking test that is run on the configured service between Router A and Router B.
<b>Action</b>	In operational mode, enter the <b>show services rpm rfc2544-benchmarking (aborted-tests   active-tests   completed-tests   summary)</b> command to display information about the results of each category or state of the RFC 2544-based benchmarking test, such as aborted tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RFC2544-Based Benchmarking Tests Overview on page 983</a></li><li>• <a href="#">Configuring an RFC 2544-Based Benchmarking Test on page 989</a></li></ul>

---

## Example: Configuring an RFC 2544-Based Benchmarking Test for UNI Direction of Ethernet Pseudowires

This example shows how to configure the benchmarking test for the user-to-network interface (UNI) direction of an Ethernet pseudowire service.

- [Requirements on page 1176](#)
- [Overview on page 1176](#)
- [Configuration on page 1177](#)
- [Verifying the Results of the Benchmarking Test for UNI Direction of an Ethernet Pseudowire Service on page 1182](#)

## Requirements

This example uses the following hardware and software components:

- An ACX Series router—f
- Junos OS Release or later

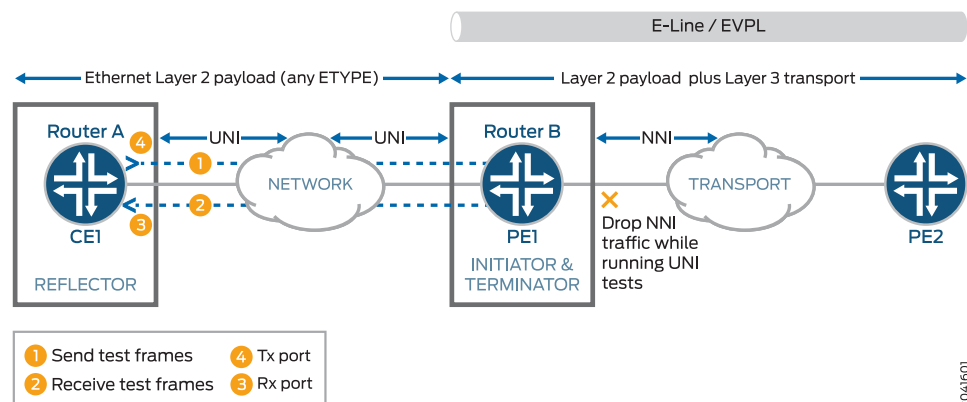
## Overview

Consider a sample topology in which a router, Router A, functions as a reflector of the test frames for an RFC 2544-based benchmarking test. The logical customer edge (CE)-facing interface and `inet` family are configured on Router A. Router A is not part of a pseudowire and therefore, a Layer 3 family configuration is required on it. Router A, which is a customer edge device CE1 is connected to Router B, which functions as a provider edge device PE1 over an Ethernet pseudowire in the UNI direction with EtherType or Layer 2 Ethernet payload. The logical interface, family, and UNI direction are configured on Router B. Router B or PE1 is connected over an Ethernet pseudowire in the NNI direction to a provider edge device at the remote site, PE2. The link between CE1 and PE1 is an Ethernet Layer 2 network and it can be configured with any EtherType value. The link between PE1 and PE2 is an Ethernet line (E-LINE) or an Ethernet Private Line (EPL) that has Layer 2 payload and Layer 3 transport sent over it. Router B or PE1 functions as an initiator and terminator of the test frames that are sent to Router A and reflected back from it.

This benchmarking test is used to compute the performance attributes in the user-to-network interface (UNI) direction of an Ethernet pseudowire service between Router A and Router B. Data traffic arriving from a network-to-network interface (NNI) toward the customer edge is ignored while the test is in progress. Packets from the CE are not sent toward the NNI because all packets are assumed to be test probes.

Figure 38 on page 1002 shows the sample topology to perform an RFC 2544 test for the UNI direction of an Ethernet pseudowire service.

**Figure 44: RFC 2544-Based Benchmarking Test for UNI Direction of an Ethernet Pseudowire**



## Configuration

In this example, you configure the benchmarking test for the UNI direction of an Ethernet pseudowire service that is enabled between two routers to detect and analyze the performance of the interconnecting routers.

- [Configuring Benchmarking Test Parameters on Router A on page 1178](#)
- [Configuring Benchmarking Test Parameters on Router B on page 1180](#)
- [Results on page 1181](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level:

#### Configuring Benchmarking Test Parameters on Router A

```
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 0 vlan-id 101
set interfaces ge-0/0/0 unit 0 family inet address 200.0.0.1/24
set services rpm rfc2544-benchmarking profiles test-profile throughput test-type
 throughput
set services rpm rfc2544-benchmarking profiles test-profile throughput packet-size 64
set services rpm rfc2544-benchmarking profiles test-profile throughput test-duration
 20m
set services rpm rfc2544-benchmarking profiles test-profile throughput bandwidth-kbps
 500
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/0.1
set services rpm rfc2544-benchmarking tests test-name test1 test-profile throughput
set services rpm rfc2544-benchmarking tests test-name test1 mode initiate,terminate
set services rpm rfc2544-benchmarking tests test-name test1 family inet
set services rpm rfc2544-benchmarking tests test-name test1 dest-address 200.0.0.2
set services rpm rfc2544-benchmarking tests test-name test1 udp-port 4001
```

#### Configuring Benchmarking Test Parameters on Router B

```
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 0 encapsulation vlan-ccc
set interfaces ge-0/0/4 unit 0 vlan-id 101
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/4.1
set services rpm rfc2544-benchmarking tests test-name test1 mode reflect
set services rpm rfc2544-benchmarking tests test-name test1 mode family ccc
set services rpm rfc2544-benchmarking tests test-name test1 direction uni
```

### Configuring Benchmarking Test Parameters on Router A

---

**Step-by-Step Procedure** The following require you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the test parameters on Router A:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:  

```
[edit]
user@host# edit interfaces
```
2. Configure the interface on which the test must be run.  

```
[edit interfaces]
user@host# edit ge-0/0/0
```
3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.  

```
[edit interfaces ge-0/0/0]
user@host# set vlan-tagging
```
4. Configure a logical unit and specify the protocol family as **inet**.  

```
[edit interfaces ge-0/0/0]
user@host# edit unit 0 family inet
```
5. Specify the address for the logical interface.  

```
[edit interfaces ge-0/0/0 unit 0 family inet]
user@host# set address 200.0.0.1/24
```
6. Configure the VLAN ID on the logical interface as 101.  

```
[edit interfaces ge-0/0/0 unit 0]
user@host# set vlan-id 101
```
7. Go to the top level of the configuration command mode.  

```
[edit interfaces ge-0/0/0 unit 0]
user@host# top
```
8. In configuration mode, go to the **[edit services]** hierarchy level.  

```
[edit]
user@host# edit services
```
9. Configure a real-time performance monitoring service (RPM) instance.  

```
[edit services]
user@host# edit rpm
```
10. Configure an RFC 2544-based benchmarking test for the RPM instance.  

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```
11. Define a name for a test profile—for example, throughput.  

```
[edit services rpm rfc2544-benchmarking]
user@host# edit profiles test-profile throughput
```

12. Configure the type of test to be performed as throughput.  

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type throughput
```
13. Specify the size of the test packet as 64 bytes.  

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type packet-size 64
```
14. Specify the period for which the test is to be performed in hours, minutes, or seconds by specifying a number followed by the letter h (for hours), m (for minutes), or s (for seconds). In this example, you configure the period as 20 minutes.  

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type test-duration 20m
```
15. Define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.  

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type bandwidth-kbps 500
```
16. Enter the **up** command to go the previous level in the configuration hierarchy.  

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# up
```
17. Enter the **up** command to go the previous level in the configuration hierarchy.  

```
[edit services rpm rfc2544-benchmarking profiles]
user@host# up
```
18. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.  

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```
19. Specify the name of the test profile—for example, throughput—to be associated with a particular test name.  

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-profile throughput
```
20. Specify the logical interface, ge-0/0/0.1, on which the RFC 2544-based benchmarking test is run.  

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface ge-0/0/0.1
```
21. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.  

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode initiate-and-terminate
```
22. Configure the address type family, **inet**, for the benchmarking test.  

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family inet
```
23. Configure the destination IPv4 address for the test packets as 200.0.0.2.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set dest-address 200.0.0.2
```

24. Specify the UDP port of the destination to be used in the UDP header for the generated frames as 4001.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set udp-port 4001
```

---

### Configuring Benchmarking Test Parameters on Router B

#### Step-by-Step Procedure

The following require you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the test parameters on Router B:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

```
[edit]
user@host# edit interfaces
```

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@host# edit ge-0/0/4
```

3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.

```
[edit interfaces ge-0/0/4]
user@host# set vlan-tagging
```

4. Configure a logical unit for the interface.

```
[edit interfaces ge-0/0/4]
user@host# edit unit 0
```

5. Specify the encapsulation for Ethernet VLAN circuits.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# set encapsulation vlan-ccc
```

6. Configure the VLAN ID as 101 on the logical interface.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# set vlan-id 101
```

7. Go to the top level of the configuration command mode.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# top
```

8. In configuration mode, go to the **[edit services]** hierarchy level.

```
[edit]
user@host# edit services
```

9. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]
user@host# edit rpm
```

10. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

11. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

12. Specify the logical interface on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface ge-0/0/4.1
```

13. Specify **reflect** as the test mode for the packets that are sent during the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```

14. Configure the address type family, **ccc**, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```

15. Specify the direction of the interface on which the test must be run, which is UNI in this example.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction uni
```

16. Start the benchmarking test on the reflector.

```
user@host> test services rpm rfc2544-benchmarking test test1 start
```

After the test is successfully completed at the initiator, you can stop the test at the reflector by entering the **test services rpm rfc2544-benchmarking test test1 stop** command.

## Results

In configuration mode, confirm your configuration on Router A and Router B by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Benchmarking Test Parameters on Router A:

```
[edit interfaces]
ge-0/0/0 {
 vlan-tagging;
 unit 0 {
 vlan-id 101;
 family inet {
 address 200.0.0.1/24;
 }
 }
}
```

```
[edit services rpm]
rfc2544-benchmarking {
 profiles {
 test-profile throughput {
 test-type throughput
 packet-size 64;
 test-duration 20m;
 bandwidth-kbps 500;
 }
 }

 tests {
 test-name test1 {
 interface ge-0/0/0.1;
 test-profile throughput;
 mode initiate,terminate;
 family inet;
 dest-address 200.0.0.2
 udp-port 4001;
 }
 }
}
```

Benchmarking Test Parameters on Router B:

```
[edit interfaces]
ge-0/0/4 {
 vlan-tagging;
 unit 0 {
 encapsulation vlan-ccc;
 vlan-id 101;
 }
}

[edit services rpm]
rfc2544-benchmarking {
 # Note, When in reflector mode, test profile is not needed
 tests {
 test-name test1 {
 interface ge-0/0/4.1;
 mode reflect;
 family ccc;
 direction uni;
 }
 }
}
```

After you have configured the device, enter the **commit** command in configuration mode.

## Verifying the Results of the Benchmarking Test for UNI Direction of an Ethernet Pseudowire Service

Examine the results of the benchmarking test that is performed on the configured service between Router A and Router B.

- [Verifying the Benchmarking Test Results on page 1183](#)



### Verifying the Benchmarking Test Results

---

<b>Purpose</b>	Verify that the necessary and desired statistical values are displayed for the benchmarking test that is run on the configured service between Router A and Router B.
<b>Action</b>	In operational mode, enter the <b>show services rpm rfc2544-benchmarking (aborted-tests   active-tests   completed-tests   summary)</b> command to display information about the results of each category or state of the RFC 2544-based benchmarking test, such as aborted tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance.
<b>Meaning</b>	The output displays the details of the benchmarking test that was performed. For more information about the <b>show services rpm rfc2544-benchmarking</b> operational command, see <b>show services rpm rfc2544-benchmarking</b> in the <a href="#">CLI Explorer</a> .
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RFC2544-Based Benchmarking Tests Overview on page 983</a></li><li>• <a href="#">Configuring an RFC 2544-Based Benchmarking Test on page 989</a></li></ul>

## Example: Configuring an RFC 2544-Based Benchmarking Test for NNI Direction of Ethernet Pseudowires

---

This example shows how to configure the benchmarking test for a network-to-network interface (NNI) direction of an Ethernet pseudowire service.

- [Requirements on page 1183](#)
- [Overview on page 1183](#)
- [Configuration on page 1184](#)
- [Verifying the Results of the Benchmarking Test for NNI Direction of an Ethernet Pseudowire Service on page 1190](#)

### Requirements

This example uses the following hardware and software components:

- An ACX Series router
- Junos OS Release or later

### Overview

Consider a sample topology in which a router, Router A, functions as an initiator and terminator of the test frames for an RFC 2544-based benchmarking test. Router A operates as a provider edge device PE1, which is connected to a customer edge device CE1 on one side and over an Ethernet pseudowire to another router Router B, which functions as a reflector to reflect back the test frames it receives from Router A. Router B operates as a provider edge device, PE2, which is the remote router located at the other side of the service provider core. The UNI direction of CE1 is connected to the NNI direction

of PE1. An MPLS tunnel connects PE1 and PE2 over the Ethernet pseudowire or the Ethernet line (E-LINE).

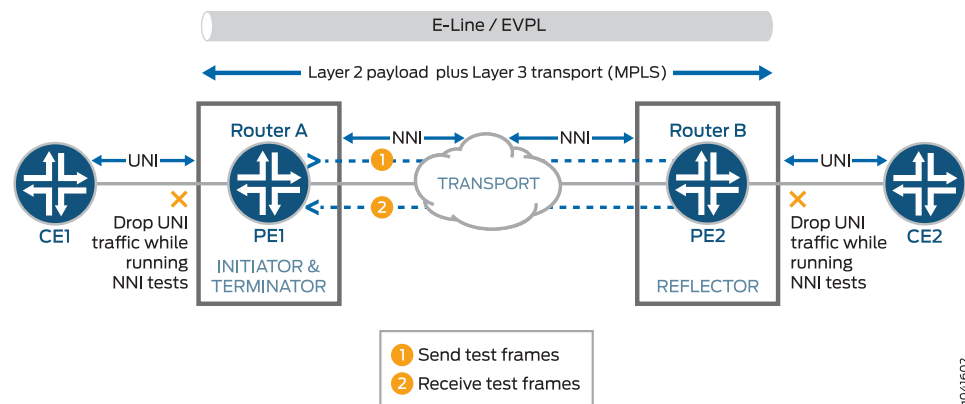


**NOTE:** When pseudowire reflection is enabled on an interface, the router does not block the ingress or egress traffic through the test interface. To block other data traffic, you must explicitly configure firewall filters.

This benchmarking test is used to compute the performance attributes in the network-to-network interface (NNI) direction of an Ethernet pseudowire service between Router A and Router B. The logical interface under test on Router A is the CE1 interface with UNI as the direction, and the logical interface under test on Router B is the CE2 interface with NNI as the direction. Data traffic arriving from UNI toward NNI is ignored while the test is in progress. Packets from NNI are not sent toward the customer edge because all packets are assumed to be test frames. The family and NNI direction are configured on routers A and B.

Figure 39 on page 1010 shows the sample topology to perform an RFC 2544 test for the NNI direction of an Ethernet pseudowire service.

**Figure 45: RFC 2544-Based Benchmarking Test for NNI Direction of an Ethernet Pseudowire**



## Configuration

In this example, you configure the benchmarking test for the NNI direction of an Ethernet pseudowire service that is enabled between two routers to detect and analyze the performance of the interconnecting routers.

- [Configuring Benchmarking Test Parameters on Router on page 1185](#)
- [Configuring Benchmarking Test Parameters on Router B on page 1187](#)
- [Results on page 1189](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level:

### Configuring Benchmarking Test Parameters on Router A

```
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 0 encapsulation vlan-ccc
set interfaces ge-0/0/0 unit 0 vlan-id 101
set services rpm rfc2544-benchmarking profiles test-profile throughput test-type
throughput
set services rpm rfc2544-benchmarking profiles test-profile throughput packet-size 64
set services rpm rfc2544-benchmarking profiles test-profile throughput test-duration 20
set services rpm rfc2544-benchmarking profiles test-profile throughput bandwidth-kbps
500
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/0.1
set services rpm rfc2544-benchmarking tests test-name test1 test-profile throughput
set services rpm rfc2544-benchmarking tests test-name test1 mode initiate,terminate
set services rpm rfc2544-benchmarking tests test-name test1 family ccc
set services rpm rfc2544-benchmarking tests test-name test1 direction nni
```

### Configuring Benchmarking Test Parameters on Router B

```
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 0 encapsulation vlan-ccc
set interfaces ge-0/0/4 unit 0 vlan-id 101
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/4.1
set services rpm rfc2544-benchmarking tests test-name test1 mode reflect
set services rpm rfc2544-benchmarking tests test-name test1 mode family ccc
set services rpm rfc2544-benchmarking tests test-name test1 direction uni
```

## Configuring Benchmarking Test Parameters on Router

### Step-by-Step Procedure

The following require you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the test parameters on Router A:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:  

```
[edit]
user@host# edit interfaces
```
2. Configure the interface on which the test must be run.  

```
[edit interfaces]
user@host# edit ge-0/0/0
```
3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.  

```
[edit interfaces ge-0/0/0]
user@host# set vlan-tagging
```
4. Configure a logical unit for the interface.

- ```
[edit interfaces ge-0/0/0]
user@host# edit unit 0
```
5. Specify the encapsulation for Ethernet VLAN circuits.

```
[edit interfaces ge-0/0/0 unit 0]
user@host# set encapsulation vlan-ccc
```
 6. Configure the VLAN ID on the logical interface.

```
[edit interfaces ge-0/0/0 unit 0]
user@host# set vlan-id 101
```
 7. Go to the top level of the configuration command mode.

```
[edit interfaces ge-0/0/0 unit 0]
user@host# top
```
 8. In configuration mode, go to the **[edit services]** hierarchy level.

```
[edit]
user@host# edit services
```
 9. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]
user@host# edit rpm
```
 10. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```
 11. Define a name for a test profile—for example, throughput.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit profiles test-profile throughput
```
 12. Configure the type of test to be performed as throughput.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type throughput
```
 13. Specify the size of the test packet as 64 bytes.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type packet-size 64
```
 14. Specify the period—for example, 20 minutes—for which the test is to be performed in hours, minutes, or seconds by specifying a number followed by the letter h (for hours), m (for minutes), or s (for seconds).

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type test-duration 20m
```
 15. Define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type bandwidth-kbps 500
```
 16. Enter the **up** command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
```

```
user@host# up
```

17. Enter the **up** command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles]  
user@host# up
```

18. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]  
user@host# edit tests test-name test1
```

19. Specify the name of the test profile—for example, throughput—to be associated with a particular test name.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set test-profile throughput
```

20. Specify the logical interface, ge-0/0/0.1, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set test-interface ge-0/0/0.1
```

21. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set mode initiate-and-terminate
```

22. Configure the address type family, ccc, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set family ccc
```

23. Specify the direction of the interface on which the test must be run, which is NNI in this example.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set direction nni
```

Configuring Benchmarking Test Parameters on Router B

Step-by-Step Procedure

The following require you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the test parameters on Router B:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

```
[edit]  
user@host# edit interfaces
```

2. Configure the interface on which the test must be run.

```
[edit interfaces]  
user@host# edit ge-0/0/4
```

3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.

```
[edit interfaces ge-0/0/4]  
user@host# set vlan-tagging
```
4. Configure a logical unit for the interface.

```
[edit interfaces ge-0/0/4]  
user@host# edit unit 0
```
5. Specify the encapsulation for Ethernet VLAN circuits.

```
[edit interfaces ge-0/0/4 unit 0]  
user@host# set encapsulation vlan-ccc
```
6. Configure the VLAN ID on the logical interface.

```
[edit interfaces ge-0/0/4 unit 0]  
user@host# set vlan-id 101
```
7. Go to the top level of the configuration command mode.

```
[edit interfaces ge-0/0/4 unit 0]  
user@host# top
```
8. In configuration mode, go to the **[edit services]** hierarchy level.

```
[edit]  
user@host# edit services
```
9. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]  
user@host# edit rpm
```
10. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]  
user@host# edit rfc2544-benchmarking
```
11. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]  
user@host# edit tests test-name test1
```
12. Specify the logical interface, ge-0/0/4.1, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set test-interface ge-0/0/4.1
```



NOTE: When pseudowire reflection is enabled on an interface, the router does not block the ingress or egress traffic through the test interface. To block other data traffic, you must explicitly configure firewall filters.

13. Specify **reflect** as the test mode for the packets that are sent during the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```

14. Configure the address type family, **ccc**, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```

15. Specify the direction of the interface on which the test must be run, which is **NNI** in this example.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction nni
```

16. Start the benchmarking test on the reflector.

```
user@host> test services rpm rfc2544-benchmarking test test1 start
```

After the test is successfully completed at the initiator, you can stop the test at the reflector by entering the **test services rpm rfc2544-benchmarking test test1 stop** command.

Results

In configuration mode, confirm your configuration on Router A and Router B by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Benchmarking Test Parameters on Router A:

```
[edit interfaces]
ge-0/0/0 {
  vlan-tagging;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 101;
  }
}

[edit services rpm]
rfc2544-benchmarking {
  profiles {
    test-profile throughput {
      test-type throughput
      packet-size 64;
      test-duration 20m;
      bandwidth-kbps 500;
    }
  }

  tests {
    test-name test1 {
      interface ge-0/0/0.1;
      test-profile throughput;
      mode initiate,terminate;
      family ccc;
      direction nni;
    }
  }
}
```

Benchmarking Test Parameters on Router B:

```
[edit interfaces]
ge-0/0/4 {
  vlan-tagging;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 101;
  }
}

[edit services rpm]
rfc2544-benchmarking {
  # Note, When in reflector mode, test profile is not needed
  tests {
    test-name test1 {
      interface ge-0/0/4.1;
      mode reflect;
      family ccc;
      direction nni;
    }
  }
}
```

After you have configured the device, enter the **commit** command in configuration mode.

Verifying the Results of the Benchmarking Test for NNI Direction of an Ethernet Pseudowire Service

Examine the results of the benchmarking test that is performed on the configured service between Router A and Router B.

- [Verifying the Benchmarking Test Results on page 1190](#)

Verifying the Benchmarking Test Results

| | |
|------------------------------|---|
| Purpose | Verify that the necessary and desired statistical values are displayed for the benchmarking test that is run on the configured service between Router A and Router B. |
| Action | In operational mode, enter the show services rpm rfc2544-benchmarking (aborted-tests active-tests completed-tests summary) command to display information about the results of each category or state of the RFC 2544-based benchmarking test, such as aborted tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance. |
| Meaning | The output displays the details of the benchmarking test that was performed. For more information about the show services rpm rfc2544-benchmarking operational command, see show services rpm rfc2544-benchmarking in the CLI Explorer . |
| Related Documentation | <ul style="list-style-type: none">• RFC2544-Based Benchmarking Tests Overview on page 983• Configuring an RFC 2544-Based Benchmarking Test on page 989 |

CHAPTER 73

Tracking Streaming Media Traffic Using Inline Video Monitoring

- [Inline Video Monitoring Overview on page 1191](#)
- [Configuring Inline Video Monitoring on page 1193](#)
- [Inline Video Monitoring Syslog Messages on page 1195](#)

Inline Video Monitoring Overview

Junos OS supports inline video monitoring using Media Delivery Index (MDI) metrics.

Inline video monitoring is available on MX Series routers using only the following MPCs:

- MPCE1
- MPCE2
- MPC-16XGE

You use the **video-monitoring** statement at the **[edit services]** hierarchy level to specify monitoring criteria for two key indicators of video traffic problems: delay factor and media loss rate (MLR), and to apply these metrics to flows on designated interfaces.

Before you use the inline video monitoring feature, ensure that you understand the following terms:

- **media delivery index**—These metrics facilitate identification of buffering needs for streaming media. Buffering must be adequate to compensate for packet jitter, measured by the MDI delay factor, and quality problems indicated by lost packets, measured by the MDI media loss rate (MLR). By performing measurements under varying load conditions, you can identify sources of significant jitter or packet loss and take appropriate action.
- **delay factor** —Delay factor is the maximum observed time difference between the arrival of media data and the drain of media data. The expected drain rate is the nominal, constant traffic rate for constant bit rate streams or the computed traffic rate of variable rate media stream packet data.

For typical stream rates of 1 megabit per second and higher, an interval of one second provides an adequate sample time. The delay factor indicates how long a data stream must be buffered (delayed) at its nominal bit rate to prevent packet loss.

The delay factor suggests the minimum size of the buffer required at the next downstream node. As a stream progresses, the variation of the delay factor indicates packet bunching or packet gaps (jitter). Greater delay factor values also indicate that more network latency is needed to deliver a stream due to the need to pre-fill a receive buffer before beginning the drain to guarantee no underflow.

When the nominal drain bit rate at a receiving node is known, the delay factor's maximum indicates the size of buffer required to accommodate packet jitter.

- **Media rate variation (MRV)**—This value is the difference between the expected packet rate and actual packet rate expressed as a percentage of the expected packet rate.
- **Media loss rate (MLR)**—This value is the number of media packets lost over a configurable time interval (*interval-duration*), where the flow packets are packets carrying streaming application information. A single IP packet can contain zero or more streaming packets. For example, an IP packet typically contains seven 188-byte MPEG transport stream packets. In this case, a single IP packet loss results in seven lost packets counted (if those seven lost packets did not include null packets). Including out-of-order packets is important, because many stream consumer-type devices do not attempt to reorder packets that are received out of order.

To configure the monitoring process, define criteria templates and apply them to the interfaces and flows you want to monitor. Monitoring templates include the following criteria:

- Duration of each measurement cycle
- Flow rate information used to establish expected flow rates
- Threshold levels for media rate variation and media loss rate that trigger desired syslog alerts

For each interface you want to monitor, you can define one or more filters to select flows for monitoring. Flows are designated as input or output flows and are uniquely identified by:

- Source IP address
- Source port
- Destination IP address
- Destination port

Junos OS supports the definition of filters for up to 256 flows, which can consist of input flows, output flows, or a combination of input and output flows. These filters provide criteria for selecting flows for monitoring. If the selection criteria consist of lists of IP addresses or ports, you could exceed the maximum number of match conditions for flows. Video monitoring selects a widely variable number of flows based on flow filters. The total number of flows that can be measured at a time depends on the specific MPC card being used, as shown in [Table 41 on page 1045](#).

When you do not define input or output flow filters for a monitored interfaces, all flows on the interface are subject to monitoring.

Table 49: MPC Flow Monitoring Capacity by Model

| MPC Model | Maximum Number of Flows Monitored Simultaneously |
|-----------|--|
| MPCE1 | 1000 |
| MPCE2 | 2000 |
| MPC-16XGE | 4000 |



NOTE: Junos OS measures both UDP flows (the default) and RTP flows. Junos OS differentiates media traffic over UDP or RTP by inspecting the first byte in the UDP payload. If the first byte of the UDP payload is 0x47 (MPEG2-TS sync byte), the traffic is treated as media traffic over UDP. Traffic is treated as media traffic over RTP if the version field is 2 and the payload type is 33 in the RTP header. When neither of these criteria are met, the packet is not considered for video monitoring.

Related Documentation

- [Configuring Inline Video Monitoring on page 1045](#)
- [show services video-monitoring mdi stats fpc-slot on page 2291](#)
- [show services video-monitoring mdi errors fpc-slot on page 2285](#)
- [show services video-monitoring mdi flows fpc-slot on page 2287](#)

Configuring Inline Video Monitoring

To configure inline video monitoring, perform the following tasks.

- [Configuring Media Delivery Indexing Criteria on page 1193](#)
- [Configuring Interface Flow Criteria on page 1195](#)

Configuring Media Delivery Indexing Criteria

To configure media delivery indexing criteria:

1. In edit mode, create a named template for video monitoring.

```
user@host# edit services video-monitoring templates template-name
```

For example,

```
user@host# edit services video-monitoring templates t1
```
2. Set the duration for sampling in seconds. Flow media delivery indexing statistics are updated at the end of this interval.

```
[edit services video-monitoring templates t1]
```

```
user@host# set interval-duration 1
```



BEST PRACTICE: If you change the interval duration when a template is being used, you cause a change in the calculated number of expected packets in an measurement interval for the template. We recommend that you do not change the interval duration for a template that is in use.

3. Set the inactivity timeout.

```
[edit services video-monitoring templates t1]
user@host# set inactivity-timeout 30
```

4. Configure either **media-rate** or **layer3-packet-rate** to establish expected flow rates used to compare to monitored flow rates.



NOTE: The media rate is the configured media bit rate for the stream. The media rate is used to establish *expected packets per second (pps)*.

The layer 3 packet rate in packets per second (pps) and is used to establish *expected bits per second (bps)*.

```
[edit services video-monitoring templates t1]
user@host# set media-rate 2972400
```

5. Set delay factor thresholds for syslog message levels.

```
[edit services video-monitoring templates t1]
user@host# set delay-factor threshold info 100
user@host# set delay-factor threshold warning 200
user@host# set delay-factor threshold critical 300
```

6. Set media loss rate thresholds for syslog message levels. You can set the threshold based on number of packets lost, or percentage of packets lost.

Or

```
[edit services video-monitoring templates t1]
user@host# set media-loss-rate threshold info percentage 5
user@host# set media-loss-rate threshold warning percentage 10
user@host# set media-loss-rate threshold critical percentage 20
```

7. Set the media rate variation thresholds for syslog message levels. The threshold is based on the ratio of the *difference* between the configured media rate and the monitored media rate to the configured media rate, expressed as a percentage.

```
[edit services video-monitoring templates t1]
user@host# set media-rate-variation threshold info 10
user@host# set media-rate-variation threshold warning 15
user@host# set media-rate-variation threshold critical 20
```

Configuring Interface Flow Criteria

To configure monitoring of flows for interfaces:

1. In edit mode, identify an interface for monitoring .

```
user@host# edit services video-monitoring interfaces interface-name
```

2. Identify input flows for monitoring. Flows are uniquely identified by source IP address, source port, destination IP address, and destination port. You can restrict flow measurement by specifying values for these identifiers. You can specify individual addresses or ports or lists of addresses and ports. If you do not specify any identifiers, all flows on the interface are monitored.

```
[edit services video-monitoring interfaces interface-name]
user@host# set input-flows input-flow-name
user@host# set input-flows input-flow-name source-address address
user@host# set input-flows input-flow-name source-port port
user@host# set input-flows input-flow-name destination-address address
user@host# set input-flows input-flow-name destination-port port
```



NOTE: You can configure a maximum of 256 flow definitions. If your flow definitions contain lists of addresses and ports, you may exceed the number of match conditions. When you exceed the limits for flows or match conditions, you receive the following constraint message when you commit:

```
'interfaces xe-0/2/2.0'
  Number of flows or Number of match condition under flows exceeded
  limit
error: configuration check-out failed
```

3. Identify output flows for monitoring, using the same options listed in Step 2.
4. Identify the template used to monitor the flows on the interface.

```
[edit services video-monitoring interfaces interface-name]
set template t1
```

Related Documentation

- [Inline Video Monitoring Overview on page 1043](#)
- [templates on page 1760](#)
- [interfaces on page 1694](#)

Inline Video Monitoring Syslog Messages

The following examples show the syslog messages produced when configured video monitoring thresholds are exceeded.

`/var/log/messages`

```
Mar 11 18:36:25 tstrtr01 fpc2 [MDI] DF: 56.71 ms, exceeded threshold for
flow(src:20.0.0.2 dst:30.0.0.2 sport:1024 dport:2048) ingressing at interface
```

```
xe-2/2/1.0 with template t1.  
Mar 11 18:36:25 tstrtr01 fpc2 [MDI] MLR : 112, exceeded threshold for flow  
(src:20.0.0.2 dst:30.0.0.2 sport:1024 dport:2048) ingressing at interface  
xe-2/2/1.0 with template t1.  
Mar 11 18:36:25 tstrtr01 fpc2 [MDI] MRV : -5.67, exceeded threshold for flow  
(src:20.0.0.2 dst:30.0.0.2 sport:1024 dport:2048) ingressing at interface  
xe-2/2/1.0 with template t1.
```

Console Messages

```
NPC2(tstrtr01 vty)# [Mar 12 01:40:58.411 LOG: Critical] [MDI] MLR : 420, exceeded  
threshold for flow (src:20.0.0.2 dst:30.0.0.2 sport:1024 dport:2048) ingressing  
at interface xe-2/2/1.0 with template t1.  
[Mar 12 01:40:58.411 LOG: Critical] [MDI] MRV : -14.89, exceeded threshold for  
flow (src:20.0.0.2 dst:30.0.0.2 sport:1024 dport:2048) ingressing at interface  
xe-2/2/1.0 with template t1.  
[Mar 12 01:40:59.412 LOG: Critical] [MDI] DF: 141.74 ms, exceeded threshold for  
flow(src:20.0.0.2 dst:30.0.0.2 sport:1024 dport:2048) ingressing at interface  
xe-2/2/1.0 with template t1.
```

Related Documentation

- [Configuring Inline Video Monitoring on page 1045](#)

PART 20

Tunnel Services

- [Overview on page 1199](#)
- [Encapsulating One Protocol Over Another Using GRE Interfaces on page 1205](#)
- [Encapsulating One IP Packet Over Another Using IP-IP Interfaces on page 1211](#)
- [Filtering Unicast Packets Through Multicast Tunnel Interfaces on page 1213](#)
- [Connecting Logical Systems Using Logical Tunnel Interfaces on page 1221](#)
- [Understanding Default PIM Tunnel Configurations on page 1237](#)
- [Facilitating VRF Table Lookup Using Virtual Loopback Tunnel Interfaces on page 1239](#)
- [Enabling a VPN to Travel Through a Non-MPLS Network Using Dynamic Tunnels on page 1245](#)

Overview

- [Tunnel Services Overview on page 1199](#)
- [Configuring Tunnel Interfaces on MX Series Routers on page 1202](#)
- [Configuring Tunnel Interfaces on T4000 Routers on page 1203](#)

Tunnel Services Overview

By encapsulating arbitrary packets inside a transport protocol, tunneling provides a private, secure path through an otherwise public network. Tunnels connect discontinuous subnetworks and enable encryption interfaces, virtual private networks (VPNs), and MPLS. If you have a Tunnel Physical Interface Card (PIC) installed in your M Series or T Series router, you can configure unicast, multicast, and logical tunnels.

You can configure two types of tunnels for VPNs: one to facilitate routing table lookups and another to facilitate VPN routing and forwarding instance (VRF) table lookups.

For information about encryption interfaces, see [“Configuring Encryption Interfaces” on page 1251](#) and the *Junos OS Administration Library for Routing Devices*. For information about VPNs, see the *Junos OS VPNs Library for Routing Devices*. For information about MPLS, see the *Junos OS MPLS Applications Library for Routing Devices*.

On SRX Series devices, Generic Routing Encapsulation (GRE) and IP-IP tunnels use internal interfaces, `gr-0/0/0` and `ip-0/0/0`, respectively. The Junos OS creates these interfaces at system bootup; they are not associated with physical interfaces.

The Juniper Networks Junos OS supports the tunnel types shown in [Table 50 on page 1199](#).

Table 50: Tunnel Interface Types

| Interface | Description |
|-----------------------|---|
| <code>gr-0/0/0</code> | <p>Configurable generic routing encapsulation (GRE) interface. GRE allows the encapsulation of one routing protocol over another routing protocol.</p> <p>Within a router, packets are routed to this internal interface, where they are first encapsulated with a GRE packet and then re-encapsulated with another protocol packet to complete the GRE. The GRE interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform GRE.</p> |

Table 50: Tunnel Interface Types (*continued*)

| Interface | Description |
|-----------------|---|
| gre | <p>Internally generated GRE interface. This interface is generated by the Junos OS to handle GRE.</p> <p>NOTE: You can configure GRE interfaces (<code>gre-x/y/z</code>) only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. This type of interface does not require a Tunnel PIC. For more information about GMPLS, see the <i>Junos OS MPLS Applications Library for Routing Devices</i> and the <i>Junos OS, Release 14.2</i>.</p> |
| ip-0/0/0 | <p>Configurable IP-over-IP encapsulation (also called IP tunneling) interface. IP tunneling allows the encapsulation of one IP packet over another IP packet.</p> <p>Packets are routed to an internal interface where they are encapsulated with an IP packet and then forwarded to the encapsulating packet's destination address. The IP-IP interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform IP tunneling.</p> |
| ipip | Internally generated IP-over-IP interface. This interface is generated by the Junos OS to handle IP-over-IP encapsulation. It is not a configurable interface. |
| lt-0/0/0 | <p>The lt interface on M Series and T Series routers supports configuration of logical systems—the capability to partition a single physical router into multiple logical devices that perform independent routing tasks.</p> <p>On SRX Series devices, the lt interface is a configurable logical tunnel interface that interconnects logical systems. See the <i>Junos OS Logical Systems Configuration Guide for Security Devices</i>.</p> |
| mt-0/0/0 | <p>Internally generated multicast tunnel interface. Multicast tunnels filter all unicast packets; if an incoming packet is not destined for a 224/8-or-greater prefix, the packet is dropped and a counter is incremented.</p> <p>Within a router, packets are routed to this internal interface for multicast filtering. The multicast tunnel interface is an internal interface only and is not associated with a physical interface. If your router has a Tunnel Services PIC, the Junos OS automatically configures one multicast tunnel interface (mt-) for each virtual private network (VPN) you configure. You do not need to configure multicast tunnel interfaces. However, you can configure properties on mt- interfaces, such as the multicast-only statement.</p> |
| mtun | Internally generated multicast tunnel interface. This interface is generated by the Junos OS to handle multicast tunnel services. It is not a configurable interface. |

Table 50: Tunnel Interface Types (*continued*)

| Interface | Description |
|-----------------|--|
| pd-0/0/0 | <p>Configurable Protocol Independent Multicast (PIM) de-encapsulation interface. In PIM sparse mode, the first-hop router encapsulates packets destined for the rendezvous point router. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree.</p> <p>Within a router, packets are routed to this internal interface for de-encapsulation. The PIM de-encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform PIM de-encapsulation.</p> <p>NOTE: On SRX Series devices, this interface type is ppd0.</p> |
| pe-0/0/0 | <p>Configurable PIM encapsulation interface. In PIM sparse mode, the first-hop router encapsulates packets destined for the rendezvous point router. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree.</p> <p>Within a router, packets are routed to this internal interface for encapsulation. The PIM encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform PIM encapsulation.</p> <p>NOTE: On SRX Series devices, this interface type is ppe0.</p> |
| pimd | Internally generated PIM de-encapsulation interface. This interface is generated by the Junos OS to handle PIM de-encapsulation. It is not a configurable interface. |
| pime | Internally generated PIM encapsulation interface. This interface is generated by the Junos OS to handle PIM encapsulation. It is not a configurable interface. |
| vt-0/0/0 | <p>Configurable virtual loopback tunnel interface. Facilitates VRF table lookup based on MPLS labels. This interface type is supported on M Series and T Series routers, but not on SRX Series devices.</p> <p>To configure a virtual loopback tunnel to facilitate VRF table lookup based on MPLS labels, you specify a virtual loopback tunnel interface name and associate it with a routing instance that belongs to a particular routing table. The packet loops back through the virtual loopback tunnel for route lookup.</p> |

Related Documentation

- [GRE Keepalive Time Overview on page 1205](#)
- [Configuring Unicast Tunnels on page 1213](#)
- [Restricting Tunnels to Multicast Traffic on page 1219](#)
- [Configuring Tunnel Interfaces on MX Series Routers on page 1202](#)
- [Configuring Tunnel Interfaces on T4000 Routers on page 1203](#)

Configuring Tunnel Interfaces on MX Series Routers

Because the MX Series routers do not support Tunnel Services PICs, you create tunnel interfaces on MX Series routers by including the following statements at the **[edit chassis]** hierarchy level:

```
[edit chassis]
fpc slot-number {
  pic number {
    tunnel-services {
      bandwidth (1g | 10g | 20g | 40g);
    }
  }
}
```

fpc slot-number is the slot number of the DPC, MPC, or MIC. On the MX80 router, the range is 0 through 1. On other MX series routers, if two SCBs are installed, the range is 0 through 11. If three SCBs are installed, the range is 0 through 5 and 7 through 11.

The **pic number** On MX80 routers, if the FPC is 0, the PIC number can only be 0. If the FPC is 1, the PIC range is 0 through 3. For all other MX series routers, the range is 0 through 3.

bandwidth (1g | 10g | 20g | 40g) is the amount of bandwidth to reserve for tunnel traffic on each Packet Forwarding Engine.



NOTE: When you use MPCs and MICs, tunnel interfaces are soft interfaces and allow as much traffic as the forwarding-path allows, so it is advantageous to setup tunnel services without artificially limiting traffic by use of the **bandwidth** option. However, you *must* specify **bandwidth** when configuring tunnel services for MX Series routers with DPCs or FPCs. The GRE key option is not supported on the tunnel interfaces for DPCs on MX960 routers.

Bandwidth rates of 20 gigabits per second and 40 gigabits per second require use of an MX Series router with the 100-Gigabit Ethernet Modular Port Concentrator (MPC) and the 100-Gigabit CFP MIC.

1g indicates that 1 gigabit per second of bandwidth is reserved for tunnel traffic.

10g indicates that 10 gigabits per second of bandwidth is reserved for tunnel traffic.

20g indicates that 20 gigabits per second of bandwidth is reserved for tunnel traffic.

40g indicates that 40 gigabits per second of bandwidth is reserved for tunnel traffic.

If you specify a bandwidth that is not compatible, tunnel services are not activated. For example, you cannot specify a bandwidth of 1 Gbps for a Packet Forwarding Engine on a 10-Gigabit Ethernet 4-port DPC.

To verify that the tunnel interfaces have been created, issue the **show interfaces terse** operational mode command. For more information, see the [CLI Explorer](#). The bandwidth

that you specify determines the port number of the tunnel interfaces that are created. When you specify a bandwidth of **1g**, the port number is always 10. When you specify any other bandwidth, the port number is always 0.



NOTE: Ingress queueing and tunnel services cannot be configured on the same MPC as it causes PFE forwarding to stop. Each feature can, however, be configured and used separately.

Related Documentation

- *Example: Configuring Tunnel Interfaces on a Gigabit Ethernet 40-Port DPC*
- *Example: Configuring Tunnel Interfaces on a 10-Gigabit Ethernet 4-Port DPC*
- *Example: Configuring Tunnel Interfaces on the MPC3E*
- *bandwidth (Tunnel Services)*
- *tunnel-services (Chassis)*
- *[edit chassis] Hierarchy Level*

Configuring Tunnel Interfaces on T4000 Routers

To create tunnel interfaces on a T4000 Core Router, include the following statements at the **[edit chassis]** hierarchy level:

```
[edit chassis]
fpc slot-number {
  pic number {
    tunnel-services {
      bandwidth bandwidth-value;
    }
  }
}
```

fpc slot-number denotes the slot number of the FPC. On the T4000 router, the range is 0 through 7.



NOTE:

- This applies only to the T4000 Type 5 FPC. If any other type of FPC is configured in this slot, this configuration is ignored and no tunnel physical interface is created.
- When you use Type 5 FPCs, the tunnel interfaces are soft interfaces and allow as much traffic as the forwarding-path allows. So, it is advantageous to setup tunnel services without artificially limiting traffic by setting the **bandwidth** statement.

pic number on the T4000 router is 0 or 1.

bandwidth *bandwidth-value* is the amount of bandwidth to reserve for the tunnel traffic on each Packet Forwarding Engine. The bandwidth value accepted includes every multiple of 10g up to 100g.

If you specify a bandwidth that is not compatible, tunnel services are not activated. For example, you cannot specify a bandwidth of 1 Gbps for a Packet Forwarding Engine on a 100-Gigabit Ethernet PIC with CFP.

To verify that the tunnel interfaces have been created, issue the **show interfaces terse** operational mode command. For more information, see the *Junos Interfaces Command Reference*.

**Related
Documentation**

- *bandwidth (Tunnel Services)*
- *tunnel-services (Chassis)*
- *[edit chassis] Hierarchy Level*

Encapsulating One Protocol Over Another Using GRE Interfaces

- [GRE Keepalive Time Overview on page 1205](#)
- [Configuring GRE Keepalive Time on page 1205](#)
- [Enabling Fragmentation on GRE Tunnels on page 1208](#)

GRE Keepalive Time Overview

Generic routing encapsulation (GRE) tunnel interfaces do not have a built-in mechanism for detecting when a tunnel is down. You can enable keepalive messages to serve as the detection mechanism.

Keepalives can be configured on the physical or on the logical interface. If configured on the physical interface, keepalives are sent on all logical interfaces that are part of the physical interface. If configured on a individual logical interface, keepalives are only sent to that logical interface. In addition to configuring a keepalive, you must configure the hold time.

Related Documentation

- [Configuring GRE Keepalive Time on page 1205](#)
- [keepalive-time on page 1797](#)
- [hold-time on page 1795](#)

Configuring GRE Keepalive Time

- [Configuring Keepalive Time and Hold time for a GRE Tunnel Interface on page 1206](#)
- [Display GRE Keepalive Time Configuration on page 1206](#)
- [Display Keepalive Time Information on a GRE Tunnel Interface on page 1207](#)

Configuring Keepalive Time and Hold time for a GRE Tunnel Interface

You can configure the keepalives on a generic routing encapsulation (GRE) tunnel interface by including both the **keepalive-time** statement and the **hold-time** statement at the **[edit protocols oam gre-tunnel interface *interface-name*]** hierarchy level.



NOTE: For proper operation of keepalives on a GRE interface, you must also include the **family inet** statement at the **[edit interfaces *interface-name* unit *unit*]** hierarchy level. If you do not include this statement, the interface is marked as down.

To configure a GRE tunnel interface:

1. Configure the GRE tunnel interface at **[edit interfaces *interface-name* unit *unit-number*]** hierarchy level, where the interface name is gr-x/y/z, and the family is set as **inet**.

```
user@host# set interfaces interface-name unit unit-number family family-name
```

2. Configure the rest of the GRE tunnel interface options as explained in *Configuring a GRE Tunnel Interface Between a PE and CE Router* or *Configuring a GRE Tunnel Interface Between PE Routers* based on requirement.

To configure keepalive time for a GRE tunnel interface:

1. Configure the Operation, Administration, and Maintenance (OAM) protocol at the **[edit protocols]** hierarchy level for the GRE tunnel interface.

```
[edit]
user@host# edit protocols oam
```

2. Configure the GRE tunnel interface option for OAM protocol.

```
[edit protocols oam]
user@host# edit gre-tunnel interface interface-name
```

3. Configure the keepalive time from 1 through 50 seconds for the GRE tunnel interface.

```
[edit protocols oam gre-tunnel interface interface-name]
user@host# set keepalive-time seconds
```

4. Configure the hold time from 5 through 250 seconds. Note that the hold time must be at least twice the keepalive time.

```
[edit protocols oam gre-tunnel interface interface-name]
user@host# set hold-time seconds
```

Display GRE Keepalive Time Configuration

Purpose Display the configured keepalive time value as 10 and hold time value as 30 on a GRE tunnel interface (for example, gr-1/1/10.1).

Action To display the configured values on the GRE tunnel interface, run the **show oam gre-tunnel** command at the **[edit protocols]** hierarchy level:

Display Keepalive Time Information on a GRE Tunnel Interface

Action To verify the current status information on a GRE tunnel interface (for example, gr-3/3/0.3), run the **show interfaces gr-3/3/0.3 terse** and **show interfaces gr-3/3/0.3 extensive** operational commands.

| | | | | | |
|------------|-------|------|--------------|--------------|--------|
| Interface | Admin | Link | Proto | Local | Remote |
| gr-3/3/0.3 | up | up | inet
mpls | 200.1.3.1/24 | |

Copyright © 2015, Juniper Networks, Inc.

**NOTE:**

When the hold time expires:

- The GRE tunnel will stay up even though the interface cannot send or receive traffic.
- The Link status will be Up and the Gre keepalives adjacency state will be Down.

Meaning The current status information of a GRE tunnel interface with keepalive time and hold time parameters is displayed as expected when the hold time expires.

Related Documentation

- [GRE Keepalive Time Overview on page 1205](#)
- [keepalive-time on page 1797](#)
- [hold-time on page 1795](#)

Enabling Fragmentation on GRE Tunnels

To enable fragmentation of IPv4 packets in generic routing encapsulation (GRE) tunnels, include the **clear-dont-fragment-bit** statement and a maximum transmission unit (MTU) setting for the tunnel as part of an existing GRE configuration at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
gr-fpc/pic/port {
  unit logical-unit-number {
    clear-dont-fragment-bit;
    ...
  }
  family inet {
    mtu 1000;
    ...
  }
}
```

This statement clears the Don't Fragment (DF) bit in the packet header, regardless of the packet size. If the packet size exceeds the tunnel MTU value, the packet is fragmented before encapsulation. The maximum MTU size configurable on the AS or Multiservices PIC is 9192 bytes.



NOTE: The **clear-dont-fragment-bit** statement is supported only on MX Series routers and all M Series routers except the M320 router.

Fragmentation is enabled only on IPv4 packets being encapsulated in IPv4-based GRE tunnels.



.....

NOTE: This configuration is supported only on GRE tunnels on AS or Multiservices interfaces. If you commit `gre-fragmentation` as the encapsulation type on a standard Tunnel PIC interface, the following console log message appears when the PIC comes online:

`gr-fpc/pic/port: does not support this encapsulation`

The Packet Forwarding Engine updates the IP identification field in the outer IP header of GRE-encapsulated packets, so that reassembly of the packets is possible after fragmentation. The previous CLI constraint check that required you to configure either the `clear-dont-fragment-bit` statement or a tunnel key with the `allow-fragmentation` statement is no longer enforced.

When you configure the `clear-dont-fragment-bit` statement on an interface with the MPLS protocol family enabled, you must specify an MTU value. This MTU value must not be greater than maximum supported value (9192).

.....

Related Documentation

- [Configuring Unicast Tunnels on page 1213](#)

CHAPTER 76

Encapsulating One IP Packet Over Another Using IP-IP Interfaces

- [Configuring IPv6-over-IPv4 Tunnels on page 1211](#)
- [Example: Configuring an IPv6-over-IPv4 Tunnel on page 1211](#)

Configuring IPv6-over-IPv4 Tunnels

If you have a Tunnel PIC installed in your M Series or T Series router, you can configure IPv6-over-IPv4 tunnels. To define a tunnel, you configure a unicast tunnel across an existing IPv4 network infrastructure. IPv6/IPv4 packets are encapsulated in IPv4 headers and sent across the IPv4 infrastructure through the configured tunnel. You manually configure configured tunnels on each end point.

On SRX Series devices, Generic Routing Encapsulation (GRE) and IP-IP tunnels use internal interfaces, `gr-0/0/0` and `ip-0/0/0`, respectively. The Junos OS creates these interfaces at system bootup; they are not associated with a physical interface.

IPv6-over-IPv4 tunnels are defined in RFC 2893, *Transition Mechanisms for IPv6 Hosts and Routers*. For information about configuring a unicast tunnel, see [“Configuring Unicast Tunnels” on page 1213](#). For an IPv6-over-IPv4 tunnel configuration example, see [“Example: Configuring an IPv6-over-IPv4 Tunnel” on page 1211](#).

Related Documentation

- [Tunnel Services Overview on page 1199](#)
- [Example: Configuring an IPv6-over-IPv4 Tunnel on page 1211](#)

Example: Configuring an IPv6-over-IPv4 Tunnel

Configure a tunnel on both sides of the connection.

| | |
|--------------------------------------|--|
| Configuration on
Router 1 | <pre>[edit] interfaces { gr-1/0/0 { unit 0 { tunnel { source 10.19.2.1; destination 10.19.3.1; } } } }</pre> |
|--------------------------------------|--|

```
        family inet6 {  
            address 2001:DB8:1:1/126;  
        }  
    }  
}
```

Configuration on Router 2 [edit]

```
interfaces {  
    gr-1/0/0 {  
        unit 0 {  
            tunnel {  
                source 10.19.3.1;  
                destination 10.19.2.1;  
            }  
            family inet6 {  
                address 2001:DB8:2:1/126;  
            }  
        }  
    }  
}
```

- Related Documentation**
- [Tunnel Services Overview on page 1199](#)
 - [Configuring IPv6-over-IPv4 Tunnels on page 1211](#)

Filtering Unicast Packets Through Multicast Tunnel Interfaces

- [Configuring Unicast Tunnels on page 1213](#)
- [Examples: Configuring Unicast Tunnels on page 1218](#)
- [Restricting Tunnels to Multicast Traffic on page 1219](#)

Configuring Unicast Tunnels

To configure a unicast tunnel, you configure a **gr-** interface (to use GRE encapsulation) or an **ip-** interface (to use IP-IP encapsulation) and include the **tunnel** and **family** statements:

```
gr-fpc/pic/port or ip-fpc/pic/port {
  unit logical-unit-number {
    copy-tos-to-outer-ip-header;
    reassemble-packets;
    tunnel {
      allow-fragmentation;
      destination destination-address;
      do-not-fragment;
      key number;
      routing-instance {
        destination routing-instance-name;
      }
      source address;
      ttl number;
    }
    family family {
      address address {
        destination address;
      }
    }
  }
}
```

You can configure these statements at the following hierarchy levels:

- **[edit interfaces]**
- **[edit logical-systems *logical-system-name* interfaces]**

You can configure multiple logical units for each GRE or IP-IP interface, and you can configure only one tunnel per unit.



NOTE: On M Series and T Series routers, you can configure the interface on a service PIC or a tunnel PIC. On MX Series routers, configure the interface on a Multiservices DPC.

Each tunnel interface must be a point-to-point interface. Point to point is the default interface connection type, so you do not need to include the **point-to-point** statement in the logical interface configuration.

You must specify the tunnel's destination and source addresses. The remaining statements are optional.



NOTE: For transit packets exiting the tunnel, forwarding path features, such as reverse path forwarding (RPF), forwarding table filtering, source class usage, destination class usage, and stateless firewall filtering, are not supported on the interfaces you configure as tunnel sources, but are supported on tunnel-pic interfaces.

However, class-of-service (CoS) information obtained from the GRE or IP-IP header is carried over the tunnel and is used by the re-entering packets. For more information, see the *Class of Service Feature Guide for Routing Devices*.

To prevent an invalid configuration, the Junos OS disallows setting the address specified by the source or destination statement at the [edit interfaces *gr-fpc/pic/port* unit *logical-unit-number* tunnel] hierarchy level to be the same as the interface's own subnet address, specified by the address statement at the [edit interfaces *gr-fpc/pic/port* unit *logical-unit-number* family *family-name*] hierarchy level.

To set the time-to-live (TTL) field that is included in the encapsulating header, include the **ttl** statement. If you explicitly configure a TTL value for the tunnel, you must configure it to be one larger than the number of hops in the tunnel. For example, if the tunnel has seven hops, you must configure a TTL value of 8.

You must configure at least one family on the logical interface. To enable MPLS over GRE tunnel interfaces, you must include the **family mpls** statement in the GRE interface configuration. In addition, you must include the appropriate statements at the [edit **protocols**] hierarchy level to enable Resource Reservation Protocol (RSVP), MPLS, and label-switched paths (LSPs) over GRE tunnels. Unicast tunnels are bidirectional.

A configured tunnel cannot go through Network Address Translation (NAT) at any point along the way to the destination. For more information, see [“Examples: Configuring Unicast Tunnels” on page 1218](#) and the *MPLS Applications Feature Guide for Routing Devices*.

For a GRE tunnel, the default is to set the ToS bits in the outer IP header to all zeros. To have the Routing Engine copy the ToS bits from the inner IP header to the outer, include

the **copy-tos-bits-to-outer-ip-header** statement. (This inner-to-outer ToS bits copying is already the default behavior for IP-IP tunnels.)

For GRE tunnel interfaces on Adaptive Services or Multiservices interfaces, you can configure additional tunnel attributes, as described in the following sections:

- [Configuring a Key Number on GRE Tunnels on page 1215](#)
- [Enabling Fragmentation on GRE Tunnels on page 1216](#)
- [Specifying an MTU Setting for the Tunnel on page 1216](#)
- [Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header on page 1217](#)
- [Configuring Packet Reassembly on page 1217](#)

Configuring a Key Number on GRE Tunnels

For Adaptive Services and Multiservices interfaces on M Series and T Series routers, you can assign a key value to identify an individual traffic flow within a GRE tunnel, as defined in RFC 2890, *Key and Sequence Number Extensions to GRE*. However, only one key is allowed for each tunnel source and destination pair.

Each IP version 4 (IPv4) packet entering the tunnel is encapsulated with the GRE tunnel key value. Each IPv4 packet exiting the tunnel is verified by the GRE tunnel key value and de-encapsulated. The Adaptive Services or Multiservices PIC drops packets that do not match the configured key value.

To assign a key value to a GRE tunnel interface, include the **key** statement:

```
key number;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* tunnel]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel]

The key number can be 0 through 4,294,967,295. You must configure the same GRE tunnel key value on tunnel endpoints.

The following example illustrates the use of the key statement in a GRE tunnel configuration:

```
interfaces {
  gr-1/2/0 {
    unit 0 {
      tunnel {
        source 10.58.255.193;
        destination 10.58.255.195;
        key 1234;
      }
    }
    ...
    family inet {
      mtu 1500;
      address 10.200.0.1/30;
    }
  }
}
```

```
    ...  
  }  
}  
}
```

Enabling Fragmentation on GRE Tunnels

For GRE tunnel interfaces on Adaptive Services and Multiservices interfaces only, you can enable fragmentation of IPv4 packets in GRE tunnels.

By default, IPv4 traffic transmitted over GRE tunnels is not fragmented. To enable fragmentation of IPv4 packets in GRE tunnels, include the **clear-dont-fragment-bit** statement:

```
clear-dont-fragment-bit;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

When you include the **clear-dont-fragment-bit** statement in the configuration, the don't-fragment (DF) bit is cleared on all packets, even packets that do not exceed the tunnel maximum transmission unit (MTU). If the packet's size exceeds the tunnel's MTU value, the packet is fragmented before encapsulation. If the packet's size does not exceed the tunnel's MTU value, the packet is not fragmented.



NOTE: The Packet Forwarding Engine updates the IP identification field in the outer IP header of GRE-encapsulated packets, so that reassembly of the packets is possible after fragmentation. The previous CLI constraint check that required you to configure either the **clear-dont-fragment-bit** statement or a tunnel key with the **allow-fragmentation** statement is no longer enforced.

You can also clear the DF bit in packets transmitted over IP Security (IPsec) tunnels. For more information, see *Enabling IPsec Packet Fragmentation*.

Specifying an MTU Setting for the Tunnel

To enable key numbers and fragmentation on GRE tunnels (as described in “[Configuring a Key Number on GRE Tunnels](#)” on page 1215 and “[Enabling Fragmentation on GRE Tunnels](#)” on page 1216), you must also specify an MTU setting for the tunnel.

To specify an MTU setting for the tunnel, include the **mtu** statement:

```
mtu bytes;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *gr-fpc/pic/port* unit *logical-unit-number* family inet]

- [edit logical-system *logical-system-name* interfaces *gr-fpc/pic/port* unit *logical-unit-number* family inet]

For more information about MTU settings, see the *Junos OS Network Interfaces Library for Routing Devices*.

Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header

Unlike IP-IP tunnels, GRE tunnels do not copy the ToS bits to the outer IP header by default. To have the Routing Engine copy the inner ToS bits to the outer IP header (which is required for some tunneled routing protocols) on packets sent by the Routing Engine, include the **copy-tos-to-outer-ip-header** statement at the logical unit hierarchy level of a GRE interface. This example copies the inner ToS bits to the outer IP header on a GRE tunnel:

```
[edit interfaces]
gr-0/0/0 {
  unit 0 {
    copy-tos-to-outer-ip-header;
    family inet;
  }
}
```

Configuring Packet Reassembly

On GRE tunnel interfaces only, you can enable reassembly of fragmented tunnel packets. To activate this capability, include the **reassemble-packets** statement:

```
reassemble-packets;
```

You can configure this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For each tunnel you configure on the interface, you can enable or disable fragmentation of GRE packets by including the **allow-fragmentation** or **do-not-fragment** statement:

```
allow-fragmentation;
do-not-fragment;
```

You can configure these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* tunnel]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel]

If you configure **allow-fragmentation** on a tunnel, it clears the DF bit in the outer IP header, enabling post fragmentation of GRE-encapsulated packets if the packet size exceeds the maximum transmission unit (MTU) value for the egress interface. By default, packets that exceed the MTU size are dropped and post fragmentation of GRE packets is disabled.



NOTE: Whenever you configure **allow-fragmentation** on a tunnel, you must also include either the **tunnel key** or the **clear-dont-fragment-bit** statement. This configuration enables the router to send affected packets to the PIC so that the correct IP header can be placed in the fragments. Otherwise, on the reassembly side some packets might be lost when fragments arrive in the PIC out of sequence at high speeds.

Starting with Junos OS Release 14.2, you can configure the generic routing encapsulation (GRE) tunnel interfaces on MX Series routers with Trio-based FPCs (MPCs) to support IP packet reassembly. The IP packet is fragmented over a GRE tunnel when the packet size exceeds the maximum transmission unit (MTU) defined for the connection. When a GRE packet passes through a link that contains an MTU less than that of the packet size, the packet is fragmented. The other endpoint of the tunnel needs to reassemble it before GRE processing starts. You can configure the GRE interfaces to reassemble the fragmented packets at the endpoint of the tunnel before they can be further processed on the network by including the **reassemble-packets** statement at the **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level. You can view the reassembly statistics by using the **show services inline ip-reassembly stastics <fpc fpc-slot | pfe pfe-slot>** command. MX Series routers can be used as the data center edge or data center interconnect devices in network environments as needed to enable interconnection of data centers across an IP backbone using GRE encapsulation. Inline IP reassembly is supported on MX80, MX240, MX480, MX960, MX2010, MX2020, and MX104 routers. The line modules compatible with the corresponding MX Series routers that support the reassembly of GRE packets are MPC1, MPC2, MPC3, MPC4, and MPC-16X10GE. Until Junos OS Release 14.1, reassembly of IP fragments received at GRE tunnels is supported only on MX Series routers with MS-DPCs.

Related Documentation

- [Tunnel Services Overview on page 1199](#)
- [Examples: Configuring Unicast Tunnels on page 1218](#)

Examples: Configuring Unicast Tunnels

Configure two unnumbered IP-IP tunnels:

```
[edit interfaces]
ip-0/3/0 {
  unit 0 {
    tunnel {
      source 192.168.4.18;
      destination 192.168.4.253;
    }
    family inet;
  }
  unit 1 {
    tunnel {
      source 192.168.4.18;
      destination 192.168.4.254;
    }
  }
}
```

```

        family inet;
    }
}

```

Configure numbered tunnel interfaces by including an address at the **[edit interfaces ip-0/3/0 unit (0 | 1) family inet]** hierarchy level:

```

[edit interfaces]
ip-0/3/0 {
  unit 0 {
    tunnel {
      source 192.168.4.18;
      destination 192.168.4.253;
    }
    family inet {
      address 10.5.5.1/30;
    }
  }
  unit 1 {
    tunnel {
      source 192.168.4.18;
      destination 192.168.4.254;
    }
    family inet {
      address 10.6.6.100/30;
    }
  }
}

```

Configure an MPLS over GRE tunnel by including the **family mpls** statement at the **[edit interfaces gr-1/2/0 unit 0]** hierarchy level:

```

[edit interfaces]
gr-1/2/0 {
  unit 0 {
    tunnel {
      source 192.168.1.1;
      destination 192.168.1.2;
    }
    family inet {
      address 10.1.1.1/30;
    }
    family mpls;
  }
}

```

- Related Documentation**
- [Tunnel Services Overview on page 1199](#)
 - [Configuring Unicast Tunnels on page 1213](#)

Restricting Tunnels to Multicast Traffic

For interfaces that carry IPv4 or IP version 6 (IPv6) traffic, you can configure a tunnel interface to allow multicast traffic only. To configure a multicast-only tunnel, include the **multicast-only** statement:

multicast-only;

You can configure this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

Multicast tunnels filter all unicast packets; if an incoming packet is not destined for a 224/8 or greater prefix, the packet is dropped and a counter is incremented.

You can configure this property on GRE, IP-IP, PIM, and multicast tunnel (**mt**) interfaces only.



NOTE: If your router has a Tunnel Services PIC, the Junos OS automatically configures one multicast tunnel interface (**mt**) for each virtual private network (VPN) you configure. You do not need to configure multicast tunnel interfaces.

**Related
Documentation**

- [Tunnel Services Overview on page 1199](#)
- [Configuring Unicast Tunnels on page 1213](#)

Connecting Logical Systems Using Logical Tunnel Interfaces

- [Configuring Logical Tunnel Interfaces on page 1221](#)
- [Example: Configuring Logical Tunnels on page 1222](#)
- [Redundant Logical Tunnels Overview on page 1224](#)
- [Configuring Redundant Logical Tunnels on page 1226](#)
- [Example: Configuring Redundant Logical Tunnels on page 1227](#)

Configuring Logical Tunnel Interfaces

Logical tunnel (**lt-**) interfaces provide quite different services depending on the host router:

- On M Series, MX Series, and T Series routers, logical tunnel interfaces allow you to connect logical systems, virtual routers, or VPN instances. M Series and T Series routers must be equipped with a Tunnel Services PIC or an Adaptive Services Module (only available on M7i routers). MX Series routers must be equipped with a Trio MPC/MIC module. For more information about connecting these applications, see the *Junos OS VPNs Library for Routing Devices*.
- On SRX Series Services Gateways, the logical tunnel interface is used to interconnect logical systems. See the *Junos OS Logical Systems Configuration Guide for Security Devices*.

For M Series, MX Series, and T Series routers, see the following section:

- [Connecting Logical Systems on page 1221](#)

Connecting Logical Systems

To connect two logical systems, you configure a logical tunnel interface on both logical systems. Then you configure a peer relationship between the logical tunnel interfaces, thus creating a point-to-point connection.

To configure a point-to-point connection between two logical systems, configure the logical tunnel interface by including the **lt-fpc/pic/port** statement:

```
lt-fpc/pic/port {
```

```
unit logical-unit-number {  
    encapsulation encapsulation;  
    peer-unit unit-number; # peering logical system unit number  
    dlcid dlcid-number;  
    family (inet | inet6 | iso | mpls);  
}  
}
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces]**
- **[edit logical-systems *logical-system-name* interfaces]**

When configuring logical tunnel interfaces, note the following:

- You can configure each logical tunnel interface with one of the following encapsulation types: Ethernet, Ethernet circuit cross-connect (CCC), Ethernet VPLS, Frame Relay, Frame Relay CCC, VLAN, VLAN CCC, or VLAN VPLS.
- You can configure the IP, IPv6, International Organization for Standardization (ISO), or MPLS protocol family.
- The peering logical interfaces must belong to the same logical tunnel interface derived from the Tunnel Services PIC or Adaptive Services Module.
- You can configure only one peer unit for each logical interface. For example, unit 0 cannot peer with both unit 1 and unit 2.
- To enable the logical tunnel interface, you must configure at least one physical interface statement.
- Logical tunnels are not supported with Adaptive Services, Multiservices, or Link Services PICs (but they are supported on the Adaptive Services Module on M7i routers, as noted above).
- On M Series routers other than the M40e router, logical tunnel interfaces require an Enhanced Flexible PIC Concentrator (FPC).
- On MX Series routers, logical tunnel interfaces require Trio MPC/MIC modules. They do not require a Tunnel Services PIC in the same system.

For more information about configuring logical systems, see the *Junos OS Routing Protocols Library for Routing Devices*.

**Related
Documentation**

- [Tunnel Services Overview on page 1199](#)
- [Example: Configuring Logical Tunnels on page 1222](#)

Example: Configuring Logical Tunnels

Configure three logical tunnels:

```
[edit interfaces]  
lt-4/2/0 {  
    description "Logical tunnel interface connects three logical systems";
```



```

}
[edit logical-systems]
lr1 {
  interfaces lt-4/2/0 {
    unit 12 {
      peer-unit 21; #Peering with lr2
      encapsulation frame-relay;
      dlci 612;
      family inet;
    }
    unit 13 {
      peer-unit 31; #Peering with lr3
      encapsulation frame-relay-ccc;
      dlci 613;
    }
  }
}
lr2 {
  interfaces lt-4/2/0 {
    unit 21 {
      peer-unit 12; #Peering with lr1
      encapsulation frame-relay-ccc;
      dlci 612;
    }
    unit 23 {
      peer-unit 32; #Peering with lr3
      encapsulation frame-relay;
      dlci 623;
    }
  }
}
lr3 {
  interfaces lt-4/2/0 {
    unit 31 {
      peer-unit 13; #Peering with lr1
      encapsulation frame-relay;
      dlci 613;
      family inet;
    }
    unit 32 {
      peer-unit 23; #Peering with lr2
      encapsulation frame-relay-ccc;
      dlci 623;
    }
  }
}

```

- Related Documentation**
- [Tunnel Services Overview on page 1199](#)
 - [Configuring Logical Tunnel Interfaces on page 1221](#)

Redundant Logical Tunnels Overview

You can connect two devices, such as an access-facing device and a core-facing device, through logical tunnels. To provide redundancy for the tunnels, you can create and configure multiple physical logical tunnels and add them to a virtual redundant logical tunnel.




NOTE: Redundant logical tunnels are supported only on MX Series routers with MPCs.

For example, in an MPLS access network, you can configure multiple pseudowires between an access node and an MX Series router with MPCs and add them to a redundant logical tunnel. You can then add multiple logical tunnels to the redundant logical tunnel.

Figure 46 on page 1224 shows a redundant logical tunnel between the access node and the MX Series router.

Figure 46: Redundant Logical Tunnels



The redundant logical tunnel has peer logical interfaces at each end, `lt0.0` and `lt0.1`. You can configure router features on these interfaces for the redundant logical tunnel and its members.

Each member logical tunnel has peer logical interfaces. In Figure 46 on page 1224, `lt-0/0/10.0` and `lt-0/0/10.1` are peers.

The MX Series router performs IP lookup in the Layer 3 VPN routing and forwarding (VRF) table on the router where the pseudowires that are grouped in logical tunnels terminate.

Redundant Logical Tunnel Configuration

In Junos Releases 13.3R1, 13.3R2, and 14.1R1 you can create up to 16 redundant logical tunnels, depending on the number of Packet Forwarding Engines and the number of

loopback interfaces on each Packet Forwarding Engine on your device. For Junos Release 13.3R3, 14.1R2, and 14.2 the valid range for device-count is from 1 to 255.

You can add up to 32 logical tunnels as members of a redundant logical tunnel.

When you add more than two members to the redundant logical tunnel, they are in active mode. The traffic is load-balanced over all the tunnel members.

When you add only two members to the redundant logical tunnel, you can configure the members in one of these ways:

- Both members in active mode
- One member in active mode and the other in backup mode

Redundant Logical Tunnel Failure Detection and Failover

A logical tunnel fails and is removed from the redundant logical tunnel group, and the backup logical tunnel becomes active due to one of these events:

- A hardware failure on the MPC module occurs.
- An MPC failure occurs due to a microkernel crash.
- The MPC module is administratively shut down and removed from the redundant logical tunnel.
- A power failure on the MPC module occurs.



NOTE: You can decrease the time it takes for failure detection and failover to occur. Configure the `enhanced-ip` statement at the `[edit chassis network-services]` hierarchy level to enable Packet Forwarding Engine liveliness detection.

Related Documentation

- [Example: Configuring Redundant Logical Tunnels on page 1227](#)
- [Pseudowire Subscriber Logical Interfaces Overview](#)
- [Configuring Logical Tunnel Interfaces on page 1221](#)
- [Configuring Redundant Logical Tunnels on page 1226](#)
- [Configuring a Pseudowire Subscriber Logical Interface Device](#)

Configuring Redundant Logical Tunnels

Use redundant logical tunnels to provide redundancy for logical tunnels between two devices, such as an access-facing device and a core-facing device.

When configuring redundant logical tunnel interfaces, note the following:

- In Junos OS Release 13.3 or later, you can configure redundant logical tunnels only on MX Series routers with MPCs.

In Junos Releases 13.3R1, 13.3R2, and 14.1R1 you can create up to 16 redundant logical tunnels, depending on the number of Packet Forwarding Engines and the number of loopback interfaces on each Packet Forwarding Engine on your device. For Junos Release 13.3R3, 14.1R2, and 14.2 the valid range for device-count is from 1 to 255. The command is shown below

set chassis redundancy-group interface-type redundant-logical-tunnel device-count *[number]*;

You can add up to 32 logical tunnels as members.

- When a logical tunnel with an existing configuration joins a redundant logical tunnel, you must configure the redundant logical tunnel with the settings from the existing configuration.
- You can add member logical tunnels to a parent logical tunnel for redundancy.
- When you add more than two logical tunnels to the redundant logical tunnel, the members are in active mode by default.
- When you add only two members, you can configure the members in one of these ways:
 - Both members in active mode
 - One member in active mode and the other in backup mode

To configure a redundant logical tunnel between two devices:

1. Create the logical tunnel and redundant logical tunnel interfaces.

```
[edit chassis]
user@host# set redundancy-group interface-type redundant-logical-tunnel
device-count count
user@host# set fpc slot-number pic number tunnel-services bandwidth 1g
```

2. Bind the member logical tunnels to the redundant logical tunnel.

```
[edit interfaces]
user@host# set interface-name redundancy-group member-interface interface-name
```

3. Configure the redundant logical tunnel interfaces.
4. Attach the redundant logical tunnel interface to a Layer 2 circuit.
5. Add the peer redundant logical tunnel interface to a Layer 3 VRF instance.
6. Configure MPLS and LDP in the pseudowires and the Layer 3 VPN.

```
[edit protocols]
user@host# set mpls no-cspf
user@host# set mpls interface all
user@host# set ldp interface all
```

7. Configure BGP in the Layer 3 VPN.
8. Configure OSPF on the core-facing interfaces and the router local loopback interface.
9. Set the policy options for BGP.
10. Set the router ID and the autonomous system (AS) number.

Related Documentation

- [Example: Configuring Redundant Logical Tunnels on page 1227](#)
- [Redundant Logical Tunnels Overview on page 1224](#)

Example: Configuring Redundant Logical Tunnels

This example shows how to configure redundant logical tunnels in an MPLS access network.

- [Requirements on page 1227](#)
- [Overview on page 1227](#)
- [Configuration on page 1228](#)
- [Verification on page 1234](#)

Requirements

In Junos OS Release 13.3 or later, you can configure redundant logical tunnels only on MX Series routers with MPCs.

Overview

When a logical tunnel with an existing configuration joins a redundant logical tunnel, you must configure the redundant logical tunnel with the settings from the existing configuration.

You can add member logical tunnels to a parent logical tunnel for redundancy.

On MX Series routers with MPCs, you can configure redundant logical tunnels as follows:

- In Junos Releases 13.3R1, 13.3R2, and 14.1R1 you can create up to 16 redundant logical tunnels, depending on the number of Packet Forwarding Engines and the number of loopback interfaces on each Packet Forwarding Engine on your device. For Junos Releases 13.3R3, 14.1R2, and 14.2 the valid range for device-count is from 1 to 255. The command is shown below

```
set chassis redundancy-group interface-type redundant-logical-tunnel device-count
[number];
```

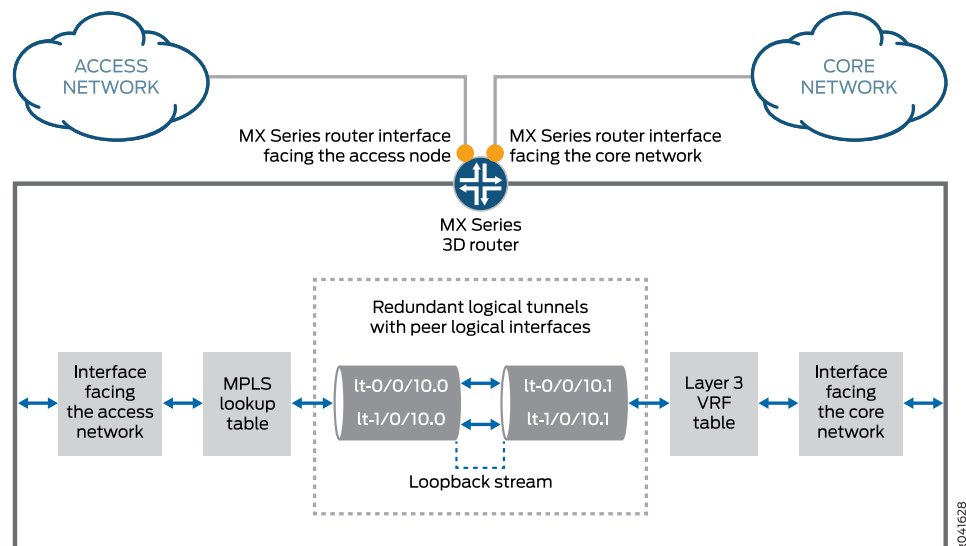
- You can add up to 32 logical tunnels as members.

- When you add more than two logical tunnels to a redundant logical tunnel, the members are in active mode by default.
- When you add only two members, you can configure the members in one of these ways:
 - Both members in active mode
 - One member in active mode and the other in backup mode

Topology

Figure 46 on page 1224 shows a redundant logical tunnel between the access node and the MX Series router in an MPLS access network.

Figure 47: Redundant Logical Tunnels



The redundant logical tunnel has peer logical interfaces at each end, `rt0.0` and `rt0.1`. You can configure router features on these interfaces for the redundant logical tunnel and its members.

Each member logical tunnel has peer logical interfaces on the access-facing and core-facing devices. In Figure 46 on page 1224, `lt-0/0/10.0` and `lt-0/0/10.1` are peers.

The MX Series router performs IP lookup in the Layer 3 VPN routing and forwarding (VRF) table on the router where the pseudowires that are grouped in logical tunnels terminate.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set chassis redundancy-group interface-type redundant-logical-tunnel device-count 4
set chassis fpc 1 pic 0 tunnel-services bandwidth 1g
```

```

set chassis fpc 1 pic 2 tunnel-services bandwidth 1g
set interfaces rlt0 redundancy-group member-interface lt-1/0/10
set interfaces rlt0 redundancy-group member-interface lt-2/0/10
set interfaces rlt0 unit 0 description "Towards Layer 2 Circuit"
set interfaces rlt0 unit 0 encapsulation vlan-ccc
set interfaces rlt0 unit 0 vlan-id 600
set interfaces rlt0 unit 0 peer-unit 1
set interfaces rlt0 unit 0 family ccc
set interfaces rlt0 unit 1 description "Towards Layer 3 VRF"
set interfaces rlt0 unit 1 encapsulation vlan
set interfaces rlt0 unit 1 vlan-id 600
set interfaces rlt0 unit 1 peer-unit 0
set interfaces rlt0 unit 1 family inet address 10.10.10.2/24
set protocols l2circuit neighbor 2.2.2.2 interface rlt0.0 virtual-circuit-id 100
set protocols l2circuit neighbor 2.2.2.2 interface rlt0.0 no-control-word
set routing-instances pe-vrf instance-type vrf
set routing-instances pe-vrf interface rlt0.1
set routing-instances pe-vrf route-distinguisher 65056:1
set routing-instances pe-vrf vrf-import VPN-A-Import
set routing-instances pe-vrf vrf-export VPN-A-Export
set routing-instances pe-vrf vrf-table-label
set routing-instances pe-vrf protocols ospf export VPN-A-Import
set routing-instances pe-vrf protocols ospf area 0.0.0.0 interface rlt0.1
set protocols mpls no-cspf
set protocols mpls interface all
set protocols ldp interface all
set protocols bgp export local-routes
set protocols bgp group internal type internal
set protocols bgp group internal local-address 3.3.3.3
set protocols bgp group internal family inet any
set protocols bgp group internal family inet-vpn unicast
set protocols bgp group internal neighbor 4.4.4.4
set protocols ospf area 0.0.0.0 interface ge-5/3/8.0
set protocols ospf area 0.0.0.0 interface ge-5/2/5.0
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set policy-options policy-statement VPN-A-Export term a then community add VPN-A
set policy-options policy-statement VPN-A-Export term a then accept
set policy-options policy-statement VPN-A-Export term b then reject
set policy-options policy-statement VPN-A-Import term a from protocol bgp
set policy-options policy-statement VPN-A-Import term a from community VPN-A
set policy-options policy-statement VPN-A-Import term a then accept
set policy-options policy-statement VPN-A-Import term b then reject
set policy-options policy-statement local-routes then accept
set policy-options community VPN-A members target:100:100
set routing-options router-id 3.3.3.3
set routing-options autonomous-system 65056

```

Step-by-Step Procedure

In this example, all the logical tunnels are in active mode.

1. Create the logical tunnel and redundant logical tunnel interfaces.

```

[edit chassis]
user@host# set redundancy-group interface-type redundant-logical-tunnel
device-count 4
user@host# set fpc 1 pic 0 tunnel-services bandwidth 1g
user@host# set fpc 1 pic 2 tunnel-services bandwidth 1g

```

2. Bind the member logical tunnels to the redundant logical tunnel.

```
[edit interfaces]
user@host# set rlt0 redundancy-group member-interface lt-1/0/10
user@host# set rlt0 redundancy-group member-interface lt-2/0/10
```

3. Configure the redundant logical tunnel interfaces.

```
[edit interfaces]
user@host# set rlt0 unit 0 description "Towards Layer 2 Circuit"
user@host# set rlt0 unit 0 encapsulation vlan-ccc
user@host# set rlt0 unit 0 vlan-id 600
user@host# set rlt0 unit 0 peer-unit 1
user@host# set rlt0 unit 0 family ccc
```

```
user@host# set rlt0 unit 1 description "Towards Layer 3 VRF"
user@host# set rlt0 unit 1 encapsulation vlan
user@host# set rlt0 unit 1 vlan-id 600
user@host# set rlt0 unit 1 peer-unit 0
user@host# set rlt0 unit 1 family inet address 10.10.10.2/24
```

4. Attach rlt0.0 to a Layer 2 circuit.

```
[edit protocols]
user@host# set l2circuit neighbor 2.2.2.2 interface rlt0.0 virtual-circuit-id 100
user@host# set l2circuit neighbor 2.2.2.2 interface rlt0.0 no-control-word
```

5. Add rlt0.1 to a Layer 3 VRF instance.

```
[edit routing-instances]
user@host# set pe-vrf instance-type vrf
user@host# set pe-vrf interface rlt0.1
user@host# set pe-vrf route-distinguisher 65056:1
user@host# set pe-vrf vrf-import VPN-A-Import
user@host# set pe-vrf vrf-export VPN-A-Export
user@host# set pe-vrf vrf-table-label
user@host# set pe-vrf protocols ospf export VPN-A-Import
user@host# set pe-vrf protocols ospf area 0.0.0.0 interface rlt0.1
```

6. Configure MPLS and LDP in the pseudowires and the Layer 3 VPN.

```
[edit protocols]
user@host# set mpls no-cspf
user@host# set mpls interface all
user@host# set ldp interface all
```

7. Configure BGP in the Layer 3 VPN.

```
[edit protocols]
user@host# set bgp export local-routes
user@host# set bgp group internal type internal
user@host# set bgp group internal local-address 3.3.3.3
user@host# set bgp group internal family inet any
user@host# set bgp group internal family inet-vpn unicast
user@host# set bgp group internal neighbor 4.4.4.4
```

8. Configure OSPF on the core-facing interfaces and the router local loopback interface.

```
[edit protocols]
user@host# set ospf area 0.0.0.0 interface ge-5/3/8.0
```



```

user@host# set ospf area 0.0.0.0 interface ge-5/2/5.0
user@host# set ospf area 0.0.0.0 interface lo0.3 passive

```

9. Set the policy options for BGP.

```

[edit policy-options]
user@host# set policy-statement VPN-A-Export term a then community add VPN-A
user@host# set policy-statement VPN-A-Export term a then accept
user@host# set policy-statement VPN-A-Export term b then reject
user@host# set policy-statement VPN-A-Import term a from protocol bgp
user@host# set policy-statement VPN-A-Import term a from community VPN-A
user@host# set policy-statement VPN-A-Import term a then accept
user@host# set policy-statement VPN-A-Import term b then reject
user@host# set policy-statement local-routes then accept
user@host# set community VPN-A members target:100:100

```

10. Set the router ID and the autonomous system (AS) number.

```

[edit routing-options]
user@host# set router-id 3.3.3.3
user@host# set autonomous-system 65056

```

Results

From configuration mode, confirm your configuration by entering the following commands:

- **show chassis**
- **show interfaces**
- **show policy-options**
- **show protocols**
- **show routing-instances**
- **show routing-options**

If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@host# show chassis
redundancy-group {
  interface-type {
    redundant-logical-tunnel {
      device-count 4;
    }
  }
}
fpc 1 {
  pic 0 {
    tunnel-services {
      bandwidth 1g;
    }
  }
}
fpc 1 {
  pic 2 {

```

```
tunnel-services {
    bandwidth 1g;
}
}

user@host# show interfaces rlt0
redundancy-group {
    member-interface lt-1/0/10;
    member-interface lt-2/0/10;
}
unit 0 {
    description "Towards Layer 2 Circuit";
    encapsulation vlan-ccc;
    vlan-id 600;
    peer-unit 1;
    family ccc;
}
unit 1 {
    description "Towards Layer 3 VRF";
    encapsulation vlan;
    vlan-id 600;
    peer-unit 0;
    family inet {
        address 10.10.10.2/24;
    }
}

user@host# show protocols l2circuit
neighbor 2.2.2.2 {
    interface rlt0.0 {
        virtual-circuit-id 100;
        no-control-word;
    }
}

user@host# show protocols
mpls {
    no-cspf;
    interface all;
}
bgp {
    export local-routes;
    group internal {
        type internal;
        local-address 3.3.3.3;
        family inet {
            any;
        }
        family inet-vpn {
            unicast;
        }
        neighbor 4.4.4.4;
    }
}
ospf {
    area 0.0.0.0 {
```

```

    interface ge-5/3/8.0;
    interface ge-5/2/5.0;
    interface lo0.3 {
        passive;
    }
}
}
ldp {
    interface all;
}
l2circuit {
    neighbor 2.2.2.2 {
        interface rlt0.0 {
            virtual-circuit-id 100;
            no-control-word;
        }
    }
}
}

user@host# routing-instances
pe-vrf {
    instance-type vrf;
    interface rlt0.1;
    route-distinguisher 65056:1;
    vrf-import VPN-A-Import;
    vrf-export VPN-A-Export;
    vrf-table-label;
    protocols {
        ospf {
            export VPN-A-Import;
            area 0.0.0.0 {
                interface rlt0.1;
            }
        }
    }
}

user@host# policy-options
policy-statement VPN-A-Export {
    term a {
        then {
            community add VPN-A;
            accept;
        }
    }
    term b {
        then reject;
    }
}
policy-statement VPN-A-Import {
    term a {
        from {
            protocol bgp;
            community VPN-A;
        }
        then accept;
    }
}

```

```
term b {  
    then reject;  
}  
}  
policy-statement local-routes {  
    then accept;  
}  
community VPN-A members target:100:100;  
  
user@host# routing-options  
router-id 3.3.3.3;  
autonomous-system 65056;
```

Verification

Confirm that the configuration is working properly.

- [Verifying the Redundant Logical Tunnel Configuration on page 1234](#)
- [Verifying the Layer 2 Circuit on page 1234](#)
- [Verifying OSPF Neighbors on page 1235](#)
- [Verifying the BGP Group on page 1235](#)
- [Verifying the BGP Routes in the Routing Table on page 1235](#)

Verifying the Redundant Logical Tunnel Configuration

Purpose Verify that the redundant logical tunnel with the child logical tunnel interfaces are created with the correct encapsulations.

Action user@host# run show interfaces terse | match rlt0

| | | | | |
|-------------|----|----|--------------|---------------|
| lt-1/0/10.0 | up | up | container--> | rlt0.0 |
| lt-1/0/10.1 | up | up | container--> | rlt0.1 |
| lt-2/0/10.0 | up | up | container--> | rlt0.0 |
| lt-2/0/10.1 | up | up | container--> | rlt0.1 |
| rlt0 | up | up | | |
| rlt0.0 | up | up | ccc | |
| rlt0.1 | up | up | inet | 10.10.10.2/24 |

Verifying the Layer 2 Circuit

Purpose Verify that the Layer 2 circuit is up.

Action user@host# run show l2circuit connections
Layer-2 Circuit Connections:

Legend for connection status (St)

| | |
|---------------------------------|--------------------------------------|
| EI -- encapsulation invalid | NP -- interface h/w not present |
| MM -- mtu mismatch | Dn -- down |
| EM -- encapsulation mismatch | VC-Dn -- Virtual circuit Down |
| CM -- control-word mismatch | Up -- operational |
| VM -- vlan id mismatch | CF -- Call admission control failure |
| OL -- no outgoing label | IB -- TDM incompatible bitrate |
| NC -- intf encaps not CCC/TCC | TM -- TDM misconfiguration |
| BK -- Backup Connection | ST -- Standby Connection |
| CB -- rcvd cell-bundle size bad | SP -- Static Pseudowire |
| LD -- local site signaled down | RS -- remote site standby |
| RD -- remote site signaled down | HS -- Hot-standby Connection |
| XX -- unknown | |

Legend for interface status

Up -- operational

Dn -- down

Neighbor: 2.2.2.2

| Interface | Type | St | Time last up | # Up trans |
|--|------|----|---------------------|------------|
| rlt0.0(vc 100) | rmt | Up | Aug 8 00:28:04 2013 | 1 |
| Remote PE: 2.2.2.2, Negotiated control-word: No | | | | |
| Incoming label: 299776, Outgoing label: 299776 | | | | |
| Negotiated PW status TLV: No | | | | |
| Local interface: rlt0.0, Status: Up, Encapsulation: VLAN | | | | |

Verifying OSPF Neighbors

Purpose Verify that routers are adjacent and able to exchange OSPF data.

Action user@host# run show ospf neighbor

| Address | Interface | State | ID | Pri | Dead |
|------------|------------|-------|---------|-----|------|
| 30.30.30.2 | ge-5/2/5.0 | Full | 4.4.4.4 | 128 | 38 |
| 20.20.20.1 | ge-5/3/8.0 | Full | 2.2.2.2 | 128 | 38 |

Verifying the BGP Group

Purpose Verify that the BGP group is created.

Action user@host# run show bgp group internal

| | | |
|--------------------------|----------------|----------------------|
| Group Type: Internal | AS: 65056 | Local AS: 65056 |
| Name: internal | Index: 0 | Flags: <Export Eval> |
| Export: [local-routes] | | |
| Holdtime: 0 | | |
| Total peers: 1 | Established: 1 | |
| 4.4.4.4+179 | | |
| inet.0: 1/6/3/0 | | |
| inet.2: 0/0/0/0 | | |
| bgp.l3vpn.0: 2/2/2/0 | | |
| pe-vrf.inet.0: 2/2/2/0 | | |

Verifying the BGP Routes in the Routing Table

Purpose Verify that the BGP routes are in the pe-vrf.inet.0 routing table.

Action user@host# run show route protocol bgp table pe-vrf.inet.0
pe-vrf.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

50.50.50.0/24 *[BGP/170] 01:18:14, localpref 100, from 4.4.4.4
 AS path: I, validation-state: unverified
 > to 30.30.30.2 via ge-5/2/5.0, Push 16

50.50.51.0/24 *[BGP/170] 01:18:14, MED 2, localpref 100, from 4.4.4.4
 AS path: I, validation-state: unverified
 > to 30.30.30.2 via ge-5/2/5.0, Push 16

- Related Documentation**
- [Configuring Redundant Logical Tunnels on page 1226](#)
 - [Redundant Logical Tunnels Overview on page 1224](#)

Understanding Default PIM Tunnel Configurations

- [Configuring PIM Tunnels on page 1237](#)

Configuring PIM Tunnels

PIM tunnels are enabled automatically on routers that have a tunnel PIC and on which you enable PIM sparse mode. You do not need to configure the tunnel interface.

PIM tunnels are unidirectional.

In PIM sparse mode, the first-hop router encapsulates packets destined for the rendezvous point (RP) router. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the RP. The RP then de-encapsulates the packets and transmits them through its multicast tree. To perform the encapsulation and de-encapsulation, the first-hop and RP routers must be equipped with Tunnel PICs.

The Junos OS creates two interfaces to handle PIM tunnels:

- **pe**—Encapsulates packets destined for the RP. This interface is present on the first-hop router.
- **pd**—De-encapsulates packets at the RP. This interface is present on the RP.



NOTE: The **pe** and **pd** interfaces do not support class-of-service (CoS) configurations.

Related Documentation

- [Tunnel Services Overview on page 1199](#)

Facilitating VRF Table Lookup Using Virtual Loopback Tunnel Interfaces

- [Configuring Virtual Loopback Tunnels for VRF Table Lookup on page 1239](#)
- [Configuring Tunnel Interfaces for Routing Table Lookup on page 1241](#)
- [Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup on page 1241](#)
- [Example: Virtual Routing and Forwarding \(VRF\) and Service Configuration on page 1242](#)

Configuring Virtual Loopback Tunnels for VRF Table Lookup

To enable egress filtering, you can either configure filtering based on the IP header, or you can configure a virtual loopback tunnel on routers equipped with a Tunnel PIC. [Table 51 on page 1239](#) describes each method.

Table 51: Methods for Configuring Egress Filtering

| Method | Interface Type | Configuration Guidelines | Comments |
|---|---|--|--|
| Filter traffic based on the IP header | Nonchannelized Point-to-Point Protocol / High Level Data Link Control (PPP/HDLC) core-facing SONET/SDH interfaces | Include the vrf-table-label statement at the [edit routing-instances instance-name] hierarchy level.

For more information, see the <i>Junos OS VPNs Library for Routing Devices</i> . | There is no restriction on customer-edge (CE) router-to-provider edge (PE) router interfaces. |
| Configure a virtual loopback tunnel on routers equipped with a Tunnel PIC | All interfaces | See the guidelines in this section. | Router must be equipped with a Tunnel PIC.

There is no restriction on the type of core-facing interface used or CE router-to-PE router interface used.

You cannot configure a virtual loopback tunnel and the vrf-table-label statement at the same time. |

You can configure a virtual loopback tunnel to facilitate VRF table lookup based on MPLS labels. You might want to enable this functionality so you can do either of the following:

- Forward traffic on a PE router to CE device interface, in a shared medium, where the CE device is a Layer 2 switch without IP capabilities (for example, a metro Ethernet switch).

The first lookup is done based on the VPN label to determine which VRF table to refer to, and the second lookup is done on the IP header to determine how to forward packets to the correct end hosts on the shared medium.

- Perform egress filtering at the egress PE router.

The first lookup on the VPN label is done to determine which VRF table to refer to, and the second lookup is done on the IP header to determine how to filter and forward packets. You can enable this functionality by configuring output filters on the VRF interfaces.

To configure a virtual loopback tunnel to facilitate VRF table lookup based on MPLS labels, you specify a virtual loopback tunnel interface name and associate it with a routing instance that belongs to a particular routing table. The packet loops back through the virtual loopback tunnel for route lookup. To specify a virtual loopback tunnel interface name, you configure the virtual loopback tunnel interface at the **[edit interfaces]** hierarchy level and include the **family inet** and **family mpls** statements:

```
vt-fpc/pic/port {  
  unit 0 {  
    family inet;  
    family mpls;  
  }  
  unit 1 {  
    family inet;  
  }  
}
```

To associate the virtual loopback tunnel with a routing instance, include the virtual loopback tunnel interface name at the **[edit routing-instances]** hierarchy level:

```
interface vt-fpc/pic/port;
```



NOTE: On virtual loopback tunnel interfaces, none of the logical interface statements except the **family** statement is supported. Note that you can configure only **inet** and **mpls** families, and you cannot configure IPv4 or IPv6 addresses on virtual loopback tunnel interfaces. Also, virtual loopback tunnel interfaces do not support class-of-service (CoS) configurations.

**Related
Documentation**

- [Tunnel Services Overview on page 1199](#)
- [Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup on page 1241](#)

Configuring Tunnel Interfaces for Routing Table Lookup

To configure tunnel interfaces to facilitate routing table lookups for VPNs, you specify a tunnel's endpoint IP addresses and associate them with a routing instance that belongs to a particular routing table. This enables the Junos OS to search in the appropriate routing table for the route prefix, because the same prefix can appear in multiple routing tables. To configure the destination VPN, include the **routing-instance** statement:

```
routing-instance {
  destination routing-instance-name;
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *gr-fpc/pic/port* unit *logical-unit-number* tunnel]
- [edit logical-systems *logical-system-name* interfaces *gr-fpc/pic/port* unit *logical-unit-number* tunnel]

This configuration indicates that the tunnel's destination address is in routing instance **routing-instance-name**. By default, the tunnel route prefixes are assumed to be in the default Internet routing table **inet.0**.



NOTE: If you configure a virtual loopback tunnel interface and the **vrf-table-label** statement on the same routing instance, the **vrf-table-label** statement takes precedence over the virtual loopback tunnel interface. For more information, see “Configuring Virtual Loopback Tunnels for VRF Table Lookup” on page 1239.

For more information about VPNs, see the *Junos OS VPNs Library for Routing Devices*.

Related Documentation

- [Tunnel Services Overview on page 1199](#)
- [destination \(Routing Instance\) on page 1789](#)

Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup

Configure a virtual loopback tunnel for VRF table lookup:

```
[edit routing-instances]
routing-instance-1 {
  instance-type vrf;
  interface vt-1/0/0.0;
  interface so-0/2/2.0;
  route-distinguisher 2:3;
  vrf-import VPN-A-import;
  vrf-export VPN-A-export;
  routing-options {
    static {
      route 10.0.0.0/8 next-hop so-0/2/2.0;
    }
  }
}
```

```
    }
  }
  routing-instance-2 {
    instance-type vrf;
    interface vt-1/0/0.1;
    interface so-0/3/2.0;
    route-distinguisher 4:5;
    vrf-import VPN-B-import;
    vrf-export VPN-B-export;
    routing-options {
      static {
        route 10.0.0.0/8 next-hop so-0/3/2.0;
      }
    }
  }
}
[edit interfaces]
vt-1/0/0 {
  unit 0 {
    family inet;
    family mpls;
  }
  unit 1 {
    family inet;
  }
}
```

**Related
Documentation**

- [Tunnel Services Overview on page 1199](#)
- [Configuring Virtual Loopback Tunnels for VRF Table Lookup on page 1239](#)

Example: Virtual Routing and Forwarding (VRF) and Service Configuration

The following example combines virtual routing and forwarding (VRF) and services configuration:

```
[edit policy-options]
policy-statement test-policy {
  term t1 {
    then reject;
  }
}
[edit routing-instances]
test {
  interface ge-0/2/0.0;
  interface sp-1/3/0.20;
  instance-type vrf;
  route-distinguisher 10.58.255.1:37;
  vrf-import test-policy;
  vrf-export test-policy;
  routing-options {
    static {
      route 0.0.0.0/0 next-table inet.0;
    }
  }
}
```

```
[edit interfaces]
ge-0/2/0 {
  unit 0 {
    family inet {
      service {
        input service-set nat-me;
        output service-set nat-me;
      }
    }
  }
}
sp-1/3/0 {
  unit 0 {
    family inet;
  }
  unit 20 {
    family inet;
    service-domain inside;
  }
  unit 21 {
    family inet;
    service-domain outside;
  }
}
[edit services]
stateful-firewall {
  rule allow-any-input {
    match-direction input;
    term t1 {
      then accept;
    }
  }
}
nat {
  pool hide-pool {
    address 10.58.16.100;
    port automatic;
  }
  rule hide-all-input {
    match-direction input;
    term t1 {
      then {
        translated {
          source-pool hide-pool;
          translation-type source napt-44;
        }
      }
    }
  }
}
service-set nat-me {
  stateful-firewall-rules allow-any-input;
  nat-rules hide-all-input;
  interface-service {
    service-interface sp-1/3/0.20;
  }
}
```

}

Enabling a VPN to Travel Through a Non-MPLS Network Using Dynamic Tunnels

- [Configuring Dynamic Tunnels on page 1245](#)

Configuring Dynamic Tunnels

A VPN that travels through a non-MPLS network requires a GRE tunnel. This tunnel can be either a static tunnel or a dynamic tunnel. A static tunnel is configured manually between two PE routers. A dynamic tunnel is configured using BGP route resolution.

When a router receives a VPN route that resolves over a BGP next hop that does not have an MPLS path, a GRE tunnel can be created dynamically, allowing the VPN traffic to be forwarded to that route. Only GRE IPv4 tunnels are supported.

To configure a dynamic tunnel between two PE routers, include the **dynamic-tunnels** statement:

```
dynamic-tunnels tunnel-name {  
    destination-networks prefix;  
    source-address address;  
}
```

You can configure this statement at the following hierarchy levels:

- [edit routing-options]
- [edit routing-instances *routing-instance-name* routing-options]
- [edit logical-systems *logical-system-name* routing-options]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options]

For more information about configuring routing options or BGP, see the *Junos OS Routing Protocols Library for Routing Devices*. For more information about VPNs, see the *Junos OS VPNs Library for Routing Devices*.

Related Documentation

- [Tunnel Services Overview on page 1199](#)

- [dynamic-tunnels on page 1792](#)

PART 21

Encryption Services

- [Overview on page 1249](#)
- [Sending Encrypted Traffic Through Tunnels on page 1251](#)
- [Configuring Redundancy in Case of Service Failure on page 1259](#)

CHAPTER 82

Overview

- [Encryption Overview on page 1249](#)
- [Configuring an ES Tunnel Interface for a Layer 3 VPN on page 1249](#)

Encryption Overview

The IP Security (IPsec) architecture provides a security suite for the IP version 4 (IPv4) and IP version 6 (IPv6) network layers. The suite provides functionality such as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. It also defines mechanisms for key generation and exchange, management of security associations, and support for digital certificates.

IPsec defines a security association (SA) and key management framework that can be used with any network layer protocol. The SA specifies what protection policy to apply to traffic between two IP-layer entities. For more information, see the *Junos OS Administration Library for Routing Devices*. The standards are defined in the following RFCs:

- RFC 2401, *Security Architecture for the Internet Protocol*
- RFC 2406, *IP Encapsulating Security Payload (ESP)*

Related Documentation

- [Configuring Encryption Interfaces on page 1251](#)
- [Configuring Filters for Traffic Transiting the ES PIC on page 1253](#)
- [Configuring ES PIC Redundancy on page 1259](#)
- [Configuring IPsec Tunnel Redundancy on page 1260](#)

Configuring an ES Tunnel Interface for a Layer 3 VPN

To configure an ES tunnel interface for a Layer 3 VPN, you need to configure an ES tunnel interface on the provider edge (PE) router and on the customer edge (CE) router. You also need to configure IPsec on the PE and CE routers. For more information about configuring an ES tunnel for a Layer 3 VPN, see the *Junos OS VPNs Library for Routing Devices*.

**Related
Documentation**

- [Encryption Overview on page 1249](#)
- [Configuring Encryption Interfaces on page 1251](#)
- [Configuring Filters for Traffic Transiting the ES PIC on page 1253](#)
- [Configuring ES PIC Redundancy on page 1259](#)
- [Configuring IPsec Tunnel Redundancy on page 1260](#)

Sending Encrypted Traffic Through Tunnels

- [Configuring Encryption Interfaces on page 1251](#)
- [Configuring Filters for Traffic Transiting the ES PIC on page 1253](#)

Configuring Encryption Interfaces

When you configure the encryption interface, you associate the configured SA with a logical interface. This configuration defines the tunnel, including the logical unit, tunnel addresses, maximum transmission unit (MTU), optional interface addresses, and the name of the IPsec SA to apply to traffic. To configure an encryption interface, include the following statements at the `[edit interfaces es-fpc/pic/port unit logical-unit-number]` hierarchy level:

```
family inet {
  ipsec-sa ipsec-sa; # name of security association to apply to packet
  address address; # local interface address inside local VPN
  destination address; # destination address inside remote VPN
}
tunnel {
  source source-address;
  destination destination-address;
}
```

The addresses configured as the tunnel source and destination are the addresses in the outer IP header of the tunnel.



NOTE: You must configure the tunnel source address locally on the router, and the tunnel destination address must be a valid address for the security gateway terminating the tunnel.

The ES Physical Interface Card (PIC) is supported on M Series and T Series routers.

The SA must be a valid tunnel-mode SA. The interface address and destination address listed are optional. The destination address allows the user to configure a static route to

encrypt traffic. If a static route uses that destination address as the next hop, traffic is forwarded through the portion of the tunnel in which encryption occurs.

Specifying the Security Association Name for Encryption Interfaces

The security association is the set of properties that defines the protocols for encrypting Internet traffic. To configure encryption interfaces, you specify the SA name associated with the interface by including the **ipsec-sa** statement at the **[edit interfaces es-fpc/pic/port unit logical-unit-number family inet]** hierarchy level:

```
ipsec-sa sa-name;
```

For information about configuring the security association, see [“Configuring Filters for Traffic Transiting the ES PIC” on page 1253](#).

Configuring the MTU for Encryption Interfaces

The protocol MTU value for encryption interfaces must always be less than the default interface MTU value of 3900 bytes; the configuration fails to commit if you select a greater value. To set the MTU value, include the **mtu** statement at the **[edit interfaces interface-name unit logical-unit-number family inet]** hierarchy level:

```
mtu bytes;
```

For more information, see the *Junos OS Network Interfaces Library for Routing Devices*.

Example: Configuring an Encryption Interface

Configure an IPsec tunnel as a logical interface on the ES PIC. The logical interface specifies the tunnel through which the encrypted traffic travels. The **ipsec-sa** statement associates the security profile with the interface.

```
[edit interfaces]
es-0/0/0 {
  unit 0 {
    tunnel {
      source 10.5.5.5; # tunnel source address
      destination 10.6.6.6; # tunnel destination address
    }
    family inet {
      ipsec-sa manual-sa1; # name of security association to apply to packet
      mtu 3800;
      address 10.1.1.8/32 { # local interface address inside local VPN
        destination 10.2.2.254; # destination address inside remote VPN
      }
    }
  }
}
```

Related Documentation

- [Encryption Overview on page 1249](#)
- [Configuring Filters for Traffic Transiting the ES PIC on page 1253](#)
- [Configuring ES PIC Redundancy on page 1259](#)
- [Configuring IPsec Tunnel Redundancy on page 1260](#)

Configuring Filters for Traffic Transiting the ES PIC

This section contains the following topics:

- [Traffic Overview on page 1253](#)
- [Configuring the Security Association on page 1254](#)
- [Configuring an Outbound Traffic Filter on page 1255](#)
- [Applying the Outbound Traffic Filter on page 1256](#)
- [Configuring an Inbound Traffic Filter on page 1256](#)
- [Applying the Inbound Traffic Filter to the Encryption Interface on page 1257](#)

Traffic Overview


Traffic configuration defines the traffic that must flow through the tunnel. You configure outbound and inbound firewall filters, which identify and direct traffic to be encrypted and confirm that decrypted traffic parameters match those defined for the given tunnel. The outbound filter is applied to the LAN or WAN interface for the incoming traffic you want to encrypt. The inbound filter is applied to the ES PIC to check the policy for traffic coming in from the remote host. Because of the complexity of configuring a router to forward packets, no automatic checking is done to ensure that the configuration is correct.



NOTE: The valid firewall filters statements for IPsec are **destination-port**, **source-port**, **protocol**, **destination-address**, and **source-address**.

In [Figure 48 on page 1253](#), Gateway A protects the network 10.1.1.0/24, and Gateway B protects the network 10.2.2.0/24. The gateways are connected by an IPsec tunnel. For more information about firewalls, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Figure 48: Example: IPsec Tunnel Connecting Security Gateways



The SA and ES interface for security Gateway A are configured as follows:

```
[edit security ipsec]
security-association manual-sa1 {
  manual {
    direction bidirectional {
      protocol esp;
      spi 2312;
      authentication {
        algorithm hmac-md5-96;
```

```
        key ascii-text 1234123412341234;
    }
    encryption {
        algorithm 3des-cbc;
        key ascii-text 123456789009876543211234;
    }
}
}
[edit interfaces es-0/1/0]
unit 0 {
    tunnel {
        source 10.5.5.5;
        destination 10.6.6.6;
    }
    family inet {
        ipsec-sa manual-sa1;
        address 10.1.1.8/32 {
            destination 10.2.2.254;
        }
    }
}
}
```

Configuring the Security Association

To configure the SA, include the **security-association** statement at the **[edit security]** hierarchy level:

```
security-association name {
    mode (tunnel | transport);
    manual {
        direction (inbound | outbound | bi-directional) {
            auxiliary-spi auxiliary-spi-value;
            spi spi-value;
            protocol (ah | esp | bundle);
            authentication {
                algorithm (hmac-md5-96 | hmac-sha1-96);
                key (ascii-text key | hexadecimal key);
            }
            encryption {
                algorithm (des-cbc | 3des-cbc);
                key (ascii-text key | hexadecimal key);
            }
        }
        dynamic {
            replay-window-size (32 | 64);
            ipsec-policy policy-name;
        }
    }
}
```

For more information about configuring an SA, see the *Junos OS Administration Library for Routing Devices*. For information about applying the SA to an interface, see [“Specifying the Security Association Name for Encryption Interfaces” on page 1252](#).

Configuring an Outbound Traffic Filter

To configure the outbound traffic filter, include the **filter** statement at the **[edit firewall]** hierarchy level:

```
filter filter-name {
  term term-name {
    from {
      match-conditions;
    }
    then {
      action;
      action-modifiers;
    }
  }
}
```

For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Example: Configuring an Outbound Traffic Filter

Firewall filters for outbound traffic direct the traffic through the desired IPsec tunnel and ensure that the tunneled traffic goes out the appropriate interface (see [Figure 48 on page 1253](#)). Here, an outbound firewall filter is created on security Gateway A; it identifies the traffic to be encrypted and adds it to the input side of the interface that carries the internal virtual private network (VPN) traffic:

```
[edit firewall]
filter ipsec-encrypt-policy-filter {
  term term1 {
    from {
      source-address { # local network
        10.1.1.0/24;
      }
      destination-address { # remote network
        10.2.2.0/24;
      }
    }
  }
  then ipsec-sa manual-sa1; # apply SA name to packet
  term default {
    then accept;
  }
}
```



NOTE: The source address, port, and protocol on the outbound traffic filter must match the destination address, port, and protocol on the inbound traffic filter. The destination address, port, and protocol on the outbound traffic filter must match the source address, port, and protocol on the inbound traffic filter.

Applying the Outbound Traffic Filter

After you have configured the outbound firewall filter, you apply it by including the **filter** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level:

```
filter {  
  input filter-name;  
}
```

Example: Applying the Outbound Traffic Filter

Apply the outbound traffic filter. The outbound filter is applied on the Fast Ethernet interface at the **[edit interfaces *fe-0/0/1* unit 0 family inet]** hierarchy level. Any packet matching the IPsec action term (**term 1**) on the input filter (**ipsec-encrypt-policy-filter**), configured on the Fast Ethernet interface, is directed to the ES PIC interface at the **[edit interfaces *es-0/1/0* unit 0 family inet]** hierarchy level. So, if a packet arrives from the source address **10.1.1.0/24** and goes to the destination address **10.2.2.0/24**, the Packet Forwarding Engine directs the packet to the ES PIC interface, which is configured with the **manual-sa1** SA. The ES PIC receives the packet, applies the **manual-sa1** SA, and sends the packet through the tunnel.

The router must have a route to the tunnel end point; add a static route if necessary.

```
[edit interfaces]  
fe-0/0/1 {  
  unit 0 {  
    family inet {  
      filter {  
        input ipsec-encrypt-policy-filter;  
      }  
      address 10.1.1.254/24;  
    }  
  }  
}
```

Configuring an Inbound Traffic Filter

To configure an inbound traffic filter, include the **filter** statement at the **[edit firewall]** hierarchy level:

```
filter filter-name {  
  term term-name {  
    from {  
      match-conditions;  
    }  
    then {  
      action;  
      action-modifiers;  
    }  
  }  
}
```

For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Example: Configuring an Inbound Traffic Filter

Configure an inbound firewall filter. This filter performs the final IPsec policy check and is created on security gateway A. The policy check ensures that only packets that match the traffic configured for this tunnel are accepted.

```
[edit firewall]
filter ipsec-decrypt-policy-filter {
  term term1 { # perform policy check
    from {
      source-address { # remote network
        10.2.2.0/24;
      }
      destination-address { # local network
        10.1.1.0/24;
      }
    }
  }
  then accept;
```

Applying the Inbound Traffic Filter to the Encryption Interface

After you create the inbound firewall filter, you can apply it to the ES PIC. To apply the filter to the ES PIC, include the **filter** statement at the **[edit interfaces es-fpc/pic/port unit logical-unit-number family inet filter]** hierarchy level:

```
filter {
  input filter;
}
```

The input filter is the name of the filter applied to received traffic. For a configuration example, see “[Example: Configuring an Inbound Traffic Filter](#)” on page 1257. For more information about firewall filters, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Example: Applying the Inbound Traffic Filter to the Encryption Interface

Apply the inbound firewall filter (**ipsec-decrypt-policy-filter**) to the decrypted packet to perform the final policy check. The IPsec **manual-sa1** SA is referenced at the **[edit interfaces es-1/2/0 unit 0 family inet]** hierarchy level and decrypts the incoming packet.

The Packet Forwarding Engine directs IPsec packets to the ES PIC. It uses the packet's security parameter index (SPI), protocol, and destination address to look up the SA configured on one of the ES interfaces. The IPsec **manual-sa1** SA is referenced at the **[edit interfaces es-1/2/0 unit 0 family inet]** hierarchy level and is used to decrypt the incoming packet. When the packets are processed (decrypted, authenticated, or both), the input firewall filter (**ipsec-decrypt-policy-filter**) is applied on the decrypted packet to perform the final policy check. **term1** defines the decrypted (and verified) traffic and performs the required policy check. For information about **term1**, see “[Example: Configuring an Inbound Traffic Filter](#)” on page 1257.



NOTE: The inbound traffic filter is applied after the ES PIC has processed the packet, so the decrypted traffic is defined as any traffic that the remote gateway is encrypting and sending to this router. IKE uses this filter to determine the policy required for a tunnel. This policy is used during the negotiation with the remote gateway to find the matching SA configuration.

```
[edit interfaces]
es-1/2/0 {
  unit 0 {
    tunnel {
      source 10.5.5.5; # tunnel source address
      destination 10.6.6.6; # tunnel destination address
    }
    family inet {
      filter {
        input ipsec-decrypt-policy-filter;
      }
      ipsec-sa manual-sa1; # SA name applied to packet
      address 10.1.1.8/32 { # local interface address inside local VPN
        destination 10.2.2.254; # destination address inside remote VPN
      }
    }
  }
}
```

**Related
Documentation**

- [Encryption Overview on page 1249](#)
- [Configuring Encryption Interfaces on page 1251](#)
- [Configuring ES PIC Redundancy on page 1259](#)
- [Configuring IPsec Tunnel Redundancy on page 1260](#)

Configuring Redundancy in Case of Service Failure

- [Configuring ES PIC Redundancy on page 1259](#)
- [Configuring IPsec Tunnel Redundancy on page 1260](#)

Configuring ES PIC Redundancy

You can configure ES PIC redundancy on M Series and T Series routers that have multiple ES PICs. With ES PIC redundancy, one ES PIC is active and another ES PIC is on standby. When the primary ES PIC has a servicing failure, the backup becomes active, inherits all the tunnels and SAs, and acts as the new next hop for IPsec traffic. Reestablishment of tunnels on the backup ES PIC does not require new Internet Key Exchange (IKE) negotiations. If the primary ES PIC comes online, it remains in standby and does not preempt the backup. To determine which PIC is currently active, use the **show ipsec redundancy** command.



NOTE: ES PIC redundancy is supported on M Series and T Series routers.

To configure an ES PIC as the backup, include the **backup-interface** statement at the **[edit interfaces fpc/pic/port es-options]** hierarchy level:

```
backup-interface es-fpc/pic/port;
```

Example: Configuring ES PIC Redundancy

After you create the inbound firewall filter, apply it to the master ES PIC. Here, the inbound firewall filter (**ipsec-decrypt-policy-filter**) is applied on the decrypted packet to perform the final policy check. The IPsec **manual-sa1** SA is referenced at the **[edit interfaces es-1/2/0 unit 0 family inet]** hierarchy level and decrypts the incoming packet. This example does not show SA and filter configuration. For information about SA and filter configuration, see the *Junos OS Administration Library for Routing Devices*, the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*, and “Example: Configuring an Inbound Traffic Filter” on page 1257.

```
[edit interfaces]
es-1/2/0 {
  es-options {
```

```

        backup-interface es-1/0/0;
    }
    unit 0 {
        tunnel {
            source 10.5.5.5;
            destination 10.6.6.6;
        }
        family inet {
            ipsec-sa manual-sa1;
            filter {
                input ipsec-decrypt-policy-filter;
            }
            address 10.1.1.8/32 {
                destination 10.2.2.254;
            }
        }
    }
}

```


Related Documentation

- [Encryption Overview on page 1249](#)
- [Configuring Encryption Interfaces on page 1251](#)
- [Configuring Filters for Traffic Transiting the ES PIC on page 1253](#)
- [Configuring IPsec Tunnel Redundancy on page 1260](#)

Configuring IPsec Tunnel Redundancy

You can configure IPsec tunnel redundancy by specifying a backup destination address. The local router sends keepalives to determine the remote site's reachability. When the peer is no longer reachable, a new tunnel is established. For up to 60 seconds during failover, traffic is dropped without notification being sent. [Figure 49 on page 1260](#) shows IPsec primary and backup tunnels.

Figure 49: IPsec Tunnel Redundancy



To configure IPsec tunnel redundancy, include the **backup-destination** statement at the **[edit interfaces unit *logical-unit-number* tunnel]** hierarchy level:

```

backup-destination address;
destination address;
source address;

```



NOTE: Tunnel redundancy is supported on M Series and T Series routers.

The primary and backup destinations must be on different routers.

The tunnels must be distinct from each other and policies must match.

For more information about tunnels, see *Tunnel Properties*.

**Related
Documentation**

- [Encryption Overview on page 1249](#)
- [Configuring Encryption Interfaces on page 1251](#)
- [Configuring Filters for Traffic Transiting the ES PIC on page 1253](#)
- [Configuring ES PIC Redundancy on page 1259](#)

PART 22

Configuration Statements and Operational Commands

- Configuration Statements on page 1265
- Operational Commands on page 1811

Configuration Statements

- [General Services Interfaces Configuration Statements on page 1265](#)
- [Adaptive Services Configuration Statements on page 1287](#)
- [Application Aware Services Configuration Statements on page 1533](#)
- [Link and Multilink Services Configuration Statements on page 1595](#)
- [Monitoring, Sampling, and Collection Services Configuration Statements on page 1618](#)
- [Tunnel and Encryption Services Configuration Statements on page 1783](#)

General Services Interfaces Configuration Statements

- [address \(Interfaces\) on page 1266](#)
- [applications \(Services ALGs\) on page 1267](#)
- [applications \(Services CoS\) on page 1267](#)
- [applications \(Services IDS\) on page 1268](#)
- [applications \(Services NAT\) on page 1268](#)
- [applications \(Services Stateful Firewall\) on page 1269](#)
- [close-timeout on page 1269](#)
- [cpu-load-threshold on page 1270](#)
- [facility-override on page 1271](#)
- [host \(Interfaces\) on page 1272](#)
- [inactivity-timeout on page 1272](#)
- [interfaces on page 1273](#)
- [log-prefix \(Interfaces\) on page 1273](#)
- [next-hop-service on page 1274](#)
- [open-timeout on page 1275](#)
- [port \(System Log Messages\) on page 1275](#)
- [rule-set \(Services Stateful Firewall\) on page 1276](#)
- [service-set \(Interfaces\) on page 1276](#)
- [service-set \(Services\) on page 1277](#)
- [services \(CoS\) on page 1279](#)

- [services \(IDS\) on page 1279](#)
- [services \(IPsec VPN\) on page 1280](#)
- [services \(Hierarchy\) on page 1280](#)
- [services \(Interfaces\) on page 1281](#)
- [services \(NAT\) on page 1281](#)
- [services \(L2TP\) on page 1282](#)
- [services \(L2TP System Logging\) on page 1282](#)
- [services \(Stateful Firewall\) on page 1283](#)
- [services \(System Logging\) on page 1284](#)
- [services-options on page 1285](#)
- [syslog \(Interfaces\) on page 1286](#)
- [tcp-tickles on page 1286](#)

address (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>address address {
 ...
}</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit (Interfaces) <i>logical-unit-number</i> family (Interfaces) <i>family</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit (Interfaces) <i>logical-unit-number</i> family (Interfaces) <i>family</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the interface address. |
| Options | <i>address</i> —Address of the interface. |
| Required Privilege Level | <i>interface</i> —To view this statement in the configuration.
<i>interface-control</i> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces.• Configuring the Address and Domain for Services Interfaces on page 45• <i>Junos OS Network Interfaces Library for Routing Devices</i> |

applications (Services ALGs)

| | |
|---------------------------------|--|
| Syntax | <code>applications { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the applications used in services. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • ALG Descriptions on page 277 • Configuring Application Sets on page 303 • Configuring Application Protocol Properties on page 303 • Examples: Configuring Application Protocols on page 321 • Verifying the Output of ALG Sessions |

applications (Services CoS)

| | |
|---------------------------------|---|
| Syntax | <code>applications [<i>application-name</i>];</code> |
| Hierarchy Level | [edit services cos rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | Define one or more applications to which the CoS services apply. |
| Options | <i>application-name</i> —Name of the target application. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Match Conditions in a CoS Rule • Configuring Match Conditions In CoS Rules on page 515 |

applications (Services IDS)

| | |
|---------------------------------|---|
| Syntax | <code>applications [<i>application-names</i>];</code> |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define one or more applications to which IDS applies. |
| Options | <i>application-name</i> —Name of the target application. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Match Conditions in IDS Rules on page 357 |

applications (Services NAT)

| | |
|---------------------------------|---|
| Syntax | <code>applications [<i>application-names</i>];</code> |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define one or more application protocols to which the NAT services apply. |
| Options | <i>application-name</i> —Name of the target application. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |

applications (Services Stateful Firewall)

| | |
|---------------------------------|---|
| Syntax | <code>applications [<i>application-names</i>];</code> |
| Hierarchy Level | [edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define one or more applications to which the stateful firewall services apply. |
| Options | <i>application-name</i> —Name of the target application. |
| Usage Guidelines | See “ Configuring Match Conditions in Stateful Firewall Rules ” on page 332. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

close-timeout

| | |
|---------------------------------|---|
| Syntax | <code>close-timeout <i>seconds</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options] |
| Release Information | Statement introduced in Junos OS Release 12.3. |
| Description | Configure the timeout period for Transmission Control Protocol (TCP) session tear-down. |
| Options | seconds —Timeout period.
Default: 1 second
Range: 2 through 300 seconds |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Default Timeout Settings for Services Interfaces on page 19 |

cpu-load-threshold

| | |
|--------------------------|--|
| Syntax | cpu-load-threshold <i>percentage</i> ; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options session-limit] |
| Release Information | Statement introduced in Release 13.2. |
| Description | Regulate the usage of CPU resources. When the CPU usage exceeds the value (percentage of the total available CPU resources) configured for cpu-load-threshold , the system reduces the rate of new sessions so that the existing sessions are not affected by low CPU availability. The CPU utilization is constantly monitored, and if the CPU usage remains in overload state—that is, above the cpu-load-threshold value configured—for a continuous period of 5 seconds, Junos OS reduces the session rate value configured at edit interfaces <i>interface-name</i> services-options session-limit rate by 10%. This is repeated until the CPU utilization comes down to the configured limit. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• services-options on page 1285 |

facility-override

| | |
|---------------------------------|--|
| Syntax | <code>facility-override <i>facility-name</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options syslog host <i>hostname</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Override the default facility for system log reporting. |
| Options | <p><i>facility-name</i>—Name of the facility that overrides the default assignment. Valid entries include:</p> <ul style="list-style-type: none"><code>authorization</code><code>daemon</code><code>ftp</code><code>kernel</code><code>local0</code> through <code>local7</code><code>user</code> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring System Logging for Services Interfaces on page 20 |

host (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <pre>host <i>hostname</i> {
 <i>services severity-level</i>;
 <i>facility-override facility-name</i>;
 <i>log-prefix prefix-value</i>;
 <i>port port-number</i>;
}</pre> |
| Hierarchy Level | [edit interfaces interface-name services-options syslog] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the hostname for the system logging utility. |
| Options | <p>hostname—Name of the system logging utility host machine. This can be the local Routing Engine or an external server address.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Applying Filters and Services to Interfaces on page 38 |

inactivity-timeout

| | |
|---------------------------------|--|
| Syntax | <pre>inactivity-timeout <i>seconds</i>;</pre> |
| Hierarchy Level | [edit interfaces interface-name services-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the inactivity timeout period for established flows. The timeout value configured in the application protocol definition overrides this value. |
| Options | <p>seconds—Timeout period.</p> <p>Default: 30 seconds</p> <p>Range: 4 through 86,400 seconds</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Default Timeout Settings for Services Interfaces on page 19 |


interfaces

| | |
|---------------------------------|---|
| Syntax | <code>interfaces { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure interfaces on the router. |
| Default | The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

log-prefix (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>log-prefix <i>prefix-value</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options <code>syslog</code> <i>host</i> <i>hostname</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Set the system logging prefix value. |
| Options | <i>prefix-value</i> —System logging prefix value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Junos OS Services Interfaces Library for Routing Devices</i> • Configuring System Logging for Services Interfaces on page 20 |

next-hop-service

| | |
|--|--|
| Syntax | <pre>next-hop-service {
 inside-service-interface <i>interface-name.unit-number</i>;
 outside-service-interface <i>interface-name.unit-number</i>;
 outside-service-interface-type local;
 service-interface-pool <i>name</i>;
}</pre> |
| Hierarchy Level | [edit services service-set service-set-name] |
| Release Information | Statement introduced before Junos OS Release 7.4.
service-interface-pool option added in Junos OS Release 9.3. |
| Description | Specify interface names or a service interface pool for the forwarding next-hop service set. You cannot specify both a service interface pool and an inside or outside interface. |
| Options | <p>inside-service-interface <i>interface-name.unit-number</i>—Name and logical unit number of the service interface associated with the service set applied inside the network.</p> <p>outside-service-interface <i>interface-name.unit-number</i>—Name and logical unit number of the service interface associated with the service set applied outside the network.</p> <p>outside-service-interface-type <i>interface-type</i>—Identifies the interface type of the service interface associated with the service set applied outside the network. For inline IP reassembly, set the interface type to local.</p> <p>service-interface-pool <i>name</i>—Name of the pool of logical interfaces configured at the [edit services service-interface-pools pool <i>pool-name</i>] hierarchy level. You can configure a service interface pool only if the service set has a PGCP rule configured. The service set cannot contain any other type of rule.</p> |
| <hr/> <div> NOTE: service-interface-pool is not applicable for IP reassembly configuration on L2TP.</div> <hr/> | |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">Configuring Service Sets to be Applied to Services Interfaces on page 31 |

open-timeout

| | |
|---------------------------------|---|
| Syntax | <code>open-timeout <i>seconds</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure a timeout period for Transmission Control Protocol (TCP) session establishment. |
| Options | <p><i>seconds</i>—Timeout period.</p> <p>Default: 5 seconds</p> <p>Range: 4 through 224 seconds</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Default Timeout Settings for Services Interfaces on page 19 |

port (System Log Messages)

| | |
|---------------------------------|---|
| Syntax | <code>port <i>port-number</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options syslog host <i>hostname</i>] |
| Release Information | Statement introduced in Junos OS Release 11.1. |
| Description | UDP port for system log messages on the host. The default port is 514. |
| Options | <i>port-number</i> —Port number for system log messages. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring System Logging for Services Interfaces on page 20 |

rule-set (Services Stateful Firewall)

| | |
|---------------------------------|---|
| Syntax | <code>rule-set <i>rule-set-name</i> {
 [<i>rule</i> <i>rule-names</i>];
}</code> |
| Hierarchy Level | [edit services stateful-firewall] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the rule set the router uses when applying this service. |
| Options | <i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set. |
| Usage Guidelines | See “ Configuring Stateful Firewall Rule Sets ” on page 335. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

service-set (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>service-set <i>service-set-name</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service (input output)],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service (input output)] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define one or more service sets to be applied to an interface. If you define multiple service sets, the router software evaluates the filters in the order in which they appear in the configuration. |
| Options | <i>service-set-name</i> —Identifies the service set. |
| Required Privilege Level | System—To view this statement in the configuration.
System-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Applying Filters and Services to Interfaces on page 38 |

service-set (Services)

```
Syntax  service-set service-set-name {
        allow-multicast;
        extension-service service-name {
            provider-specific-rules-configuration;
        }
        (ids-rules rule-name | ids-rule-sets rule-set-name);
        interface-service {
            service-interface interface-name;
        }
        ipsec-vpn-options {
            anti-replay-window-size bits;
            clear-dont-fragment-bit;
            ike-access-profile profile-name;
            local-gateway address;
            no-anti-replay;
            passive-mode-tunneling;
            trusted-ca [ ca-profile-names ];
            tunnel-mtu bytes;
        }
        ip-reassembly-rules rule-name;
        (ipsec-vpn-rules rule-name | ipsec-vpn-rule-sets rule-set-name);
        max-flows number;
        max-drop-flows {
            ingress ingress-flows;
            egress egress-flows;
        }
        nat-options {
            land-attack-check (ip-only | ip-port);
            max-sessions-per-subscriber session-number;
            stateful-nat64 {
                clear-dont-fragment-bit;
            }
        }
        (nat-rules rule-name | nat-rule-sets rule-set-name);
        next-hop-service {
            inside-service-interface interface-name.unit-number;
            outside-service-interface interface-name.unit-number;
            outside-service-interface-type local;
            service-interface-pool name;
        }
        (pgcp-rules rule-name | pgcp-rule-sets rule-set-name);
        (ptsp-rules rule-name | ptsp-rule-sets rule-set-name);
        service-set-options {
            bypass-traffic-on-exceeding-flow-limits;
            bypass-traffic-on-pic-failure;
            enable-asymmetric-traffic-processing;
            routing-engine-services;
            support-uni-directional-traffic;
        }
        snmp-trap-thresholds {
            flows high high-threshold | low low-threshold;
            nat-address-port high-threshold | low low-threshold;
        }
    }
```

```
    }  
  }  
  software-options {  
    dslite-ipv6-prefix-length dslite-ipv6-prefix-length;  
  }  
  (software-rules rule-name | software-rule-sets rule-set-name);  
  (stateful-firewall-rules rule-name | stateful-firewall-rule-sets rule-set-name);  
  syslog {  
    host hostname {  
      class {  
        alg-logs;  
        ids-logs;  
        nat-logs;  
        packet-logs;  
        pcp-logs;  
        session-logs <open | close>;  
        stateful-firewall-logs ;  
      }  
      services severity-level;  
      facility-override facility-name;  
      interface-service prefix-value;  
    }  
  }  
}
```

Hierarchy Level [edit services]

Release Information Statement introduced before Junos OS Release 7.4.
pgcp-rules and **pgcp-rule-sets** options added in Junos OS Release 8.4.
server-set-options option added in Junos OS Release 10.1.
ptsp-rules and **ptsp-rule-sets** options added in Junos OS Release 10.2.
software-rules and **clear-rule-sets** options added in Junos OS Release 10.4.
software-options option added in Junos OS Release 14.1.

Description Define the service set.

Options ***service-set-name***—Name of the service set.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Service Set Properties*

services (CoS)

| | |
|---------------------------------|---|
| Syntax | <code>services cos { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | Define the service rules to be applied to traffic. |
| Options | <code>cos</code> —Identifier for the class-of-service set of rules statements. |
| Required Privilege Level | <code>interface</code> —To view this statement in the configuration.
<code>interface-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring CoS Rules on page 514 |

services (IDS)

| | |
|---------------------------------|---|
| Syntax | <code>services ids { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the service rules to be applied to traffic. |
| Options | <code>ids</code> —Identifies the IDS set of rules statements. |
| Required Privilege Level | <code>interface</code> —To view this statement in the configuration.
<code>interface-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IDS Rules on page 355 |

services (IPsec VPN)

| | |
|---------------------------------|---|
| Syntax | <code>services ipsec-vpn { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the service rules to be applied to traffic. |
| Options | <code>ipsec-vpn</code> —IPsec set of rules statements. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Security Associations on page 385 |

services (Hierarchy)

| | |
|---------------------------------|---|
| Syntax | <code>services { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the service rules to be applied to traffic. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Service Set Properties</i> |

services (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>services severity-level;</code> |
| Hierarchy Level | [edit interfaces interface-name services-options syslog host hostname] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the system logging severity level. |
| Options | <p>severity-level—Assigns a severity level to the facility. Valid entries include:</p> <ul style="list-style-type: none"> • alert—Conditions that should be corrected immediately. • any—Matches any level. • critical—Critical conditions. • emergency—Panic conditions. • error—Error conditions. • info—Informational messages. • notice—Conditions that require special handling. • warning—Warning messages. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring System Logging for Services Interfaces on page 20 |

services (NAT)

| | |
|---------------------------------|--|
| Syntax | <code>services nat { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the service rules to be applied to traffic. |
| Options | nat —Identifies the NAT set of rules statements. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |

services (L2TP)

| | |
|---------------------------------|---|
| Syntax | <code>services l2tp { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the service properties to be applied to traffic. |
| Options | l2tp —Identifies the L2TP set of services statements. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• L2TP Services Configuration Overview on page 638 |

services (L2TP System Logging)

| | |
|---------------------------------|--|
| Syntax | <code>services <i>severity-level</i>;</code> |
| Hierarchy Level | [edit services l2tp tunnel-group <i>group-name</i> syslog host <i>hostname</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the system logging severity level. |
| Options | <i>severity-level</i> —Assigns a severity level to the facility. Valid entries include: <ul style="list-style-type: none">• alert—Conditions that should be corrected immediately.• any—Matches any level.• critical—Critical conditions.• emergency—Panic conditions.• error—Error conditions.• info—Informational messages.• notice—Conditions that require special handling.• warning—Warning messages. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring System Logging of L2TP Tunnel Activity on page 644 |

services (Stateful Firewall)

| | |
|---------------------------------|---|
| Syntax | <code>services stateful-firewall { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 10.4. |
| Description | Define the service rules to be applied to traffic. |
| Options | stateful-firewall —Identifies the stateful firewall set of rules statements. |
| Usage Guidelines | See <i>Junos Network Secure</i> . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

services (System Logging)

| | |
|---------------------------------|---|
| Syntax | <code>services severity-level;</code> |
| Hierarchy Level | [edit services service-set service-set-name syslog host hostname] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the severity level for system logging messages. |
| Options | <p>severity-level—Assigns a severity level to the facility. Valid entries are:</p> <ul style="list-style-type: none">• alert—Conditions that should be corrected immediately.• any—Matches any level.• critical—Critical conditions.• emergency—Panic conditions.• error—Error conditions.• info—Informational messages.• notice—Conditions that require special handling.• warning—Warning messages. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring System Logging for Service Sets on page 47 |

services-options

Syntax

```

services-options {
  cgn-pic;
  close-timeout
  fragment-limit
  disable-global-timeout-override;
  ignore-errors <alg> <tcp>;
  inactivity-non-tcp-timeout seconds;
  inactivity-tcp-timeout seconds;
  inactivity-timeout seconds
  open-timeout seconds;
  pba-interim-logging-interval seconds;
  reassembly-timeout
  session-limit {
    maximum number;
    rate new-sessions-per-second;
    cpu-load-threshold percentage;
  }
  session-timeout seconds;
  jflow-log {
    message-rate-limit messages-per-second;
  }
  syslog {
    host hostname {
      facility-override facility-name;
      log-prefix prefix-value;
      port port-number;
      services severity-level;
    }
    message-rate-limit messages-per-second;
  }
  tcp-tickles tcp-tickles;
  trio-flow-offload minimum-bytes minimum-bytes;
}

```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the service options to be applied on an interface.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- *Interface Properties*.

syslog (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <pre>syslog {
 host <i>hostname</i> {
 services <i>severity-level</i>;
 facility-override <i>facility-name</i>;
 log-prefix <i>prefix-value</i>;
 port <i>port-number</i>;
 }
 message-rate-limit <i>messages-per-second</i>;
}</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure generation of system log messages for the service set. System log information is passed to the kernel for logging in the <code>/var/log</code> directory. Any values configured in the service set definition override these values. |
| Options | The remaining statements are described separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring System Logging for Services Interfaces on page 20 |

tcp-tickles

| | |
|---------------------------------|---|
| Syntax | <pre>tcp-tickles <i>tcp-tickles</i>;</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Define the maximum number of keep-alive messages sent before a TCP session is allowed to timeout. |
| Options | <i>tcp-tickles</i> —Number of keep-alive messages.
Range: 0 through 30
Default: 4 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Default Timeout Settings for Services Interfaces on page 19 |

Adaptive Services Configuration Statements

- [\[edit services application-identification\] Hierarchy Level](#) on page 1295
- [IPsec Hierarchy Level](#) on page 1297
- [adaptive-services-pics](#) on page 1300
- [address \(Interfaces\)](#) on page 1301
- [address \(Services NAT Pool\)](#) on page 1301
- [address-allocation](#) on page 1302
- [address-range](#) on page 1302
- [aggregation](#) on page 1303
- [allow-ip-options](#) on page 1304
- [allow-multicast](#) on page 1305
- [allow-overlapping-nat-pools](#) on page 1305
- [anti-replay-window-size \(Services IPsec VPN\)](#) on page 1306
- [anti-replay-window-size \(Services Service Set\)](#) on page 1307
- [app-mapping-timeout](#) on page 1308
- [application](#) on page 1309
- [application-protocol](#) on page 1310
- [application-profile](#) on page 1312
- [application-set](#) on page 1313
- [application-sets \(Services CoS\)](#) on page 1313
- [application-sets \(Services IDS\)](#) on page 1314
- [application-sets \(Services NAT\)](#) on page 1314
- [application-sets \(Services Stateful Firewall\)](#) on page 1315
- [applications \(Services ALGs\)](#) on page 1315
- [applications \(Services CoS\)](#) on page 1316
- [applications \(Services IDS\)](#) on page 1316
- [applications \(Services NAT\)](#) on page 1317
- [applications \(Services Stateful Firewall\)](#) on page 1317
- [authentication \(Services IPsec VPN\)](#) on page 1318
- [authentication-algorithm \(Services IKE\)](#) on page 1319
- [authentication-algorithm \(Services IPsec\)](#) on page 1320
- [authentication-method \(Services IPsec VPN\)](#) on page 1321
- [auxiliary-spi \(Services IPsec VPN\)](#) on page 1322
- [backup-remote-gateway](#) on page 1322
- [bundle](#) on page 1323
- [by-destination](#) on page 1323

- [by-pair](#) on page 1324
- [by-source](#) on page 1325
- [bypass-traffic-on-exceeding-flow-limits](#) on page 1325
- [bypass-traffic-on-pic-failure](#) on page 1326
- [cgn-pic](#) on page 1326
- [cisco-interoperability](#) on page 1327
- [class](#) on page 1328
- [clear-dont-fragment-bit \(Interfaces GRE Tunnels\)](#) on page 1329
- [clear-dont-fragment-bit \(Services IPsec VPN\)](#) on page 1329
- [clear-dont-fragment-bit \(Services Service Set\)](#) on page 1330
- [clear-ike-sas-on-pic-restart](#) on page 1330
- [clear-ipsec-sas-on-pic-restart](#) on page 1331
- [compression](#) on page 1331
- [compression-device \(Interfaces\)](#) on page 1332
- [copy-dont-fragment-bit \(Services IPsec VPN\)](#) on page 1332
- [copy-dont-fragment-bit \(Services Set\)](#) on page 1333
- [data \(FTP\)](#) on page 1333
- [dead-peer-detection \(Services IPsec VPN\)](#) on page 1334
- [description \(Services IPsec VPN\)](#) on page 1334
- [destination-address \(Services CoS\)](#) on page 1335
- [destination-address \(Services IDS\)](#) on page 1335
- [destination-address \(Services IPsec VPN\)](#) on page 1336
- [destination-address \(Services NAT\)](#) on page 1336
- [destination-address \(Services Stateful Firewall\)](#) on page 1337
- [destination-address-range \(Services IDS\)](#) on page 1337
- [destination-address-range \(Services NAT\)](#) on page 1338
- [destination-address-range \(Services Stateful Firewall\)](#) on page 1338
- [destination-pool](#) on page 1339
- [destination-port](#) on page 1340
- [destination-port range](#) on page 1341
- [destination-prefix \(Services IDS\)](#) on page 1341
- [destination-prefix \(Services NAT\)](#) on page 1342
- [destination-prefix-ipv6](#) on page 1342
- [destination-prefix-list \(Services CoS\)](#) on page 1343
- [destination-prefix-list \(Services IDS\)](#) on page 1343
- [destination-prefix-list \(Services NAT\)](#) on page 1344
- [destination-prefix-list \(Services Stateful Firewall\)](#) on page 1344

- [destined-port](#) on page 1345
- [deterministic-port-block-allocation](#) on page 1346
- [dh-group](#) on page 1347
- [dial-options](#) on page 1348
- [direction](#) on page 1349
- [dns-alg-pool](#) on page 1349
- [dns-alg-prefix](#) on page 1350
- [drop-member-traffic](#) (Aggregated Multiservices) on page 1350
- [ds-lite](#) on page 1351
- [dscp](#) on page 1352
- [dynamic](#) on page 1352
- [ecmp-alb](#) on page 1353
- [ei-mapping-timeout](#) on page 1354
- [eif-flow-limit](#) on page 1354
- [enable-rejoin](#) (aggregated Multiservices) on page 1355
- [encapsulation](#) on page 1356
- [encryption](#) on page 1357
- [encryption-algorithm](#) (Services IPsec VPN) on page 1358
- [establish-tunnels](#) on page 1359
- [f-max-period](#) on page 1359
- [facility-override](#) (Service Sets) on page 1360
- [facility-override](#) (System Log Reporting) on page 1361
- [family](#) (Aggregated Multiservices) on page 1361
- [family](#) (Interfaces) on page 1362
- [family](#) (Voice Services) on page 1363
- [force-entry](#) on page 1364
- [forwarding-class](#) (Services CoS) on page 1364
- [forwarding-class](#) (Services CoS Fragmentation Properties) on page 1365
- [fragment-threshold](#) (Forwarding Class Maps) on page 1366
- [fragment-threshold](#) (Interfaces LSQ) on page 1367
- [fragmentation-map](#) on page 1367
- [fragmentation-maps](#) on page 1368
- [from](#) (Services CoS) on page 1369
- [from](#) (Services IDS) on page 1370
- [from](#) (Services IPsec VPN) on page 1371
- [from](#) (Services HCM) on page 1371
- [from](#) (Services NAT) on page 1372

- [from \(Services Stateful Firewall\) on page 1373](#)
- [ftp \(Services CoS\) on page 1374](#)
- [hello-interval on page 1374](#)
- [hide-avps on page 1375](#)
- [high-availability-options \(aggregated Multiservices\) on page 1376](#)
- [host \(L2TP\) on page 1377](#)
- [host \(service-set\) on page 1378](#)
- [host \(Services HCM\) on page 1379](#)
- [hot-standby on page 1379](#)
- [icmp-code on page 1380](#)
- [icmp-type on page 1380](#)
- [ids-rules on page 1381](#)
- [ignore-entry on page 1381](#)
- [ike on page 1382](#)
- [ike-access-profile on page 1383](#)
- [inactivity-timeout on page 1383](#)
- [initiate-dead-peer-detection on page 1384](#)
- [input \(Interfaces\) on page 1384](#)
- [interface-service on page 1385](#)
- [interfaces \(Aggregated Multiservices\) on page 1386](#)
- [interfaces \(Voice Services\) on page 1387](#)
- [interval on page 1387](#)
- [ipsec \(Services IPsec VPN\) on page 1388](#)
- [ipsec-inside-interface on page 1388](#)
- [ipsec-vpn-options on page 1389](#)
- [ipsec-vpn-rules on page 1389](#)
- [ipv6-multicast-interfaces on page 1390](#)
- [l2tp-access-profile on page 1390](#)
- [learn-sip-register on page 1391](#)
- [lifetime-seconds \(Services IPsec VPN\) on page 1391](#)
- [link-layer-overhead on page 1392](#)
- [load-balance on page 1392](#)
- [load-balancing-options \(Aggregated Multiservices\) on page 1393](#)
- [local-certificate \(Services IPsec VPN\) on page 1394](#)
- [local-gateway \(IPSec\) on page 1395](#)
- [local-gateway \(L2TP LNS\) on page 1395](#)
- [local-id on page 1396](#)

- [log-prefix \(L2TP\)](#) on page 1396
- [log-prefix \(Services\)](#) on page 1397
- [logging \(Services\)](#) on page 1397
- [logging \(Services IDS\)](#) on page 1398
- [lsq-failure-options](#) on page 1398
- [manual](#) on page 1399
- [many-to-one \(Aggregated Multiservices\)](#) on page 1400
- [mapping-refresh](#) on page 1401
- [mapping-timeout](#) on page 1402
- [match-direction \(Services CoS\)](#) on page 1402
- [match-direction \(Services IDS\)](#) on page 1403
- [match-direction \(Services IPsec VPN\)](#) on page 1403
- [match-direction \(Services NAT\)](#) on page 1404
- [match-direction \(Services Stateful Firewall\)](#) on page 1404
- [max-drop-flows](#) on page 1405
- [max-flows](#) on page 1406
- [maximum-contexts](#) on page 1407
- [maximum-send-window](#) on page 1407
- [member-failure-options \(Aggregated Multiservices\)](#) on page 1408
- [member-interface \(Aggregated Multiservices\)](#) on page 1410
- [message-rate-limit](#) on page 1411
- [mlfr-uni-nni-bundles-inline](#) on page 1412
- [mode \(Services IPsec VPN\)](#) on page 1413
- [mss](#) on page 1413
- [multi-link-layer-2-inline](#) on page 1414
- [multilink-class](#) on page 1414
- [multilink-max-classes](#) on page 1415
- [nat-options](#) on page 1415
- [nat-rules](#) on page 1416
- [next-hop-service](#) on page 1417
- [no-anti-replay \(Services IPsec VPN\)](#) on page 1418
- [no-anti-replay \(Services Service Set\)](#) on page 1418
- [no-fragmentation](#) on page 1419
- [no-ipsec-tunnel-in-traceroute](#) on page 1419
- [no-per-unit-scheduler](#) on page 1420
- [no-termination-request](#) on page 1420
- [no-translation](#) on page 1421

- [output](#) on page 1421
- [overload-pool](#) on page 1422
- [overload-prefix](#) on page 1422
- [passive-mode-tunneling](#) on page 1423
- [per-unit-scheduler](#) on page 1424
- [perfect-forward-secrecy](#) (Services IPsec VPN) on page 1425
- [pgcp-rules](#) on page 1426
- [policy](#) (Services IKE) on page 1427
- [policy](#) (Services IPsec VPN) on page 1428
- [pool](#) on page 1429
- [port](#) (Services NAT) on page 1430
- [port](#) (Services Voice) on page 1432
- [port](#) (System Log Messages) on page 1432
- [port-forwarding](#) on page 1433
- [port-forwarding-mappings](#) on page 1433
- [ports-per-session](#) on page 1434
- [post-service-filter](#) on page 1434
- [ppp-access-profile](#) on page 1435
- [pre-shared-key](#) (Services IKE) on page 1435
- [preserve-interface](#) on page 1436
- [primary](#) (Adaptive Services Interfaces) on page 1436
- [primary](#) (Link Services IQ PIC Interfaces) on page 1437
- [proposal](#) (Services IKE) on page 1437
- [proposal](#) (Services IPsec VPN) on page 1438
- [proposals](#) on page 1438
- [protocol](#) (Applications) on page 1439
- [protocol](#) (IPSec) on page 1440
- [ptsp-rules](#) on page 1440
- [queues](#) on page 1441
- [receive-window](#) on page 1441
- [redistribute-all-traffic](#) (Aggregated Multiservices) on page 1442
- [redundancy-options](#) (Adaptive Services Interfaces) on page 1443
- [redundancy-options](#) (Link Services IQ PIC Interfaces) on page 1443
- [\(reflexive | reverse\)](#) on page 1444
- [rejoin-timeout](#) (Aggregated Multiservices) on page 1445
- [remote-gateway](#) on page 1445
- [remote-id](#) on page 1446

- [request-url](#) on page 1446
- [retransmit-interval \(Services\)](#) on page 1447
- [rpc-program-number](#) on page 1447
- [rtp](#) on page 1448
- [rule \(Services CoS\)](#) on page 1449
- [rule \(Services IDS\)](#) on page 1450
- [rule \(Services IPsec VPN\)](#) on page 1452
- [rule \(Services NAT\)](#) on page 1454
- [rule \(Services Stateful Firewall\)](#) on page 1455
- [rule \(Softwire\)](#) on page 1456
- [rule-set \(Services CoS\)](#) on page 1456
- [rule-set \(Services IDS\)](#) on page 1457
- [rule-set \(Services IPsec VPN\)](#) on page 1457
- [rule-set \(Services NAT\)](#) on page 1458
- [rule-set \(Services Stateful Firewall\)](#) on page 1458
- [rule-set \(Softwire\)](#) on page 1459
- [secondary \(Adaptive Services Interfaces\)](#) on page 1459
- [secondary \(Link Services IQ PIC Interfaces\)](#) on page 1460
- [secure-nat-mapping](#) on page 1460
- [secured-port-block-allocation](#) on page 1461
- [server \(pcp\)](#) on page 1463
- [service](#) on page 1464
- [service-domain](#) on page 1465
- [service-filter \(Interfaces\)](#) on page 1465
- [service-interface \(Adaptive Services Interfaces\)](#) on page 1466
- [service-interface \(L2TP Processing\)](#) on page 1466
- [service-set \(Interfaces\)](#) on page 1467
- [service-set \(Services\)](#) on page 1468
- [service-set-options](#) on page 1470
- [session-limit](#) on page 1471
- [set-dont-fragment-bit \(Services Set\)](#) on page 1472
- [set-dont-fragment-bit \(Services IPsec VPN\)](#) on page 1472
- [sip-call-hold-timeout](#) on page 1473
- [sip](#) on page 1474
- [snmp-command](#) on page 1474
- [softwire-concentrator](#) on page 1475
- [softwire-options](#) on page 1476

- [software-rules](#) on page 1476
- [source-address \(Service Sets\)](#) on page 1477
- [source-address \(Services CoS\)](#) on page 1477
- [source-address \(Services IDS\)](#) on page 1478
- [source-address \(Services IPsec VPN\)](#) on page 1478
- [source-address \(Services NAT\)](#) on page 1479
- [source-address \(Services Stateful Firewall\)](#) on page 1479
- [source-address-range \(Services IDS\)](#) on page 1480
- [source-address-range \(Services NAT\)](#) on page 1480
- [source-address-range \(Services Stateful Firewall\)](#) on page 1481
- [source-pool](#) on page 1481
- [source-port](#) on page 1482
- [source-prefix \(Services IDS\)](#) on page 1482
- [source-prefix \(Services NAT\)](#) on page 1483
- [source-prefix-ipv6](#) on page 1483
- [source-prefix-list \(Services CoS\)](#) on page 1484
- [source-prefix-list \(Services IDS\)](#) on page 1484
- [source-prefix-list \(Services NAT\)](#) on page 1485
- [source-prefix-list \(Services Stateful Firewall\)](#) on page 1485
- [spi](#) on page 1486
- [stateful-firewall-rules](#) on page 1486
- [syslog \(Services CoS\)](#) on page 1487
- [syslog \(Services IDS\)](#) on page 1487
- [syslog \(Services IPsec VPN\)](#) on page 1488
- [syslog \(Services L2TP\)](#) on page 1488
- [syslog \(Services NAT\)](#) on page 1489
- [syslog \(Services Service Set\)](#) on page 1490
- [syslog \(Services Stateful Firewall\)](#) on page 1491
- [syn-cookie](#) on page 1492
- [tcp-mss](#) on page 1493
- [term \(Services CoS\)](#) on page 1494
- [term \(Services IDS\)](#) on page 1495
- [term \(Services IPsec VPN\)](#) on page 1497
- [term \(Services HCM\)](#) on page 1498
- [term \(Services NAT\)](#) on page 1499
- [term \(Services Stateful Firewall\)](#) on page 1500
- [then \(Services CoS\)](#) on page 1501

- [then \(Services HCM\) on page 1501](#)
- [then \(Services IDS\) on page 1502](#)
- [then \(Services IPsec VPN\) on page 1503](#)
- [then \(Services NAT\) on page 1504](#)
- [then \(Services Stateful Firewall\) on page 1505](#)
- [threshold \(Services IPsec\) on page 1506](#)
- [threshold \(Services Logging and SYN-Cookie Defenses\) on page 1506](#)
- [traceoptions \(Security PKI\) on page 1507](#)
- [traceoptions \(Services IPsec VPN\) on page 1509](#)
- [traceoptions \(Services L2TP\) on page 1511](#)
- [traceoptions \(Services Logging\) on page 1515](#)
- [translated on page 1517](#)
- [trigger-link-failure on page 1517](#)
- [translated-port on page 1518](#)
- [translation-type on page 1519](#)
- [trusted-ca on page 1520](#)
- [ttl-threshold on page 1521](#)
- [tunnel-group on page 1522](#)
- [tunnel-mtu \(Services IPsec VPN\) on page 1523](#)
- [tunnel-mtu \(Services Service Set\) on page 1524](#)
- [tunnel-timeout on page 1525](#)
- [url on page 1525](#)
- [url-list on page 1526](#)
- [url-rule on page 1526](#)
- [url-rule-set on page 1527](#)
- [unit \(Aggregated Multiservices\) on page 1527](#)
- [unit \(Interfaces\) on page 1528](#)
- [unit \(Voice Services\) on page 1529](#)
- [uuid on page 1530](#)
- [v6rd on page 1531](#)
- [version \(IKE\) on page 1532](#)
- [video \(Application Profile\) on page 1532](#)
- [voice \(Application Profile\) on page 1533](#)
- [warm-standby on page 1533](#)

[edit services application-identification] Hierarchy Level

To configure application identification services (APPID), include the **application-identification** statement at the **[edit services]** hierarchy level:

```
[edit services]
application-identification {
  application application-name {
    disable;
    idle-timeout seconds;
    index number;
    session-timeout seconds;
    type type;
    type-of-service service-type;
    port-mapping {
      port-range {
        tcp (port | range);
        udp (port | range);
      }
      disable;
    }
  }
}
application-group group-name {
  application-groups {
    name [application-group-name];
  }
  applications {
    name [application-name];
  }
  index number;
  disable;
}
application-system-cache-timeout seconds;
enable-heuristics
max-checked-bytes bytes;
min-checked-bytes bytes;
nested-application
nested-application-settings
no-application-identification;
no-application-system-cache;
no-clear-application-system-cache;
no-protocol-method;
no-signature-based;
profile profile-name {
  [ rule-set rule-set-name ];
}
rule rule-name {
  disable;
  address address-name {
    destination {
      ip address</prefix-length>;
      port-range {
        tcp [ ports-and-port-ranges ];
        udp [ ports-and-port-ranges ];
      }
    }
  }
  source {
    ip address</prefix-length>;
    port-range {
      tcp [ ports-and-port-ranges ];
      udp [ ports-and-port-ranges ];
    }
  }
}
```

```

    }
  }
  order number;
}
application application-name;
}
rule-set rule-set-name {
  rule application-rule-name;
}
signature-method-all-ports
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag;
  no-remote-trace;
}
}
[edit services]
hcm {
  url-rule url-rule-name {
    term term-num {
      from {
        url-list url-list-name ;
        url url_identifier {
          host hostname ;
          request-url page-name ;
        }
      }
    }
    then {
      discard;
      accept;
      count;
      log-request;
    }
  }
}
url-rule-set url-rule-set-name {
  url-rule rule1 ;
  url-rule rule2 ;
}
}
}

```

Related Documentation

- [Defining an Application Identification on page 674](#)
- [Application Identification for Nested Applications on page 681](#)
- [Configuring Global APPID Properties on page 683](#)

IPsec Hierarchy Level

To configure IP Security (IPsec) services, include the following statements at the **[edit services ipsec-vpn]** hierarchy level:

```
[edit services ipsec-vpn]
```

```
clear-ike-sas-on-pic-restart;
clear-ipsec-sas-on-pic-restart;
establish-tunnels (immediately | on-traffic);
ike {
  proposal proposal-name {
    authentication-algorithm (md5 | sha1 | sha-256 | sha-384);
    authentication-method ( pre-shared-keys | rsa-signatures);
    description description;
    dh-group (group1 | group2 | group5 | group14 | group19 | group20);
    encryption-algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
    lifetime-seconds seconds;
  }
  policy policy-name {
    description description;
    local-certificate identifier;
    local-id (ipv4_addr ipv4-address | ipv6_addr ipv6-address | key-id identifier | fqdn fqdn);
    version (1 | 2);
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposals [ proposal-names ];
    remote-id {
      any-remote-id;
      ipv4_addr [ values ];
      ipv6_addr [ values ];
      key_id [ values ];
      fqdn [ values ];
    }
  }
}
ipsec {
  proposal proposal-name {
    authentication-algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
    description description;
    encryption-algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
    lifetime-seconds seconds;
    protocol (ah | esp | bundle);
  }
  policy policy-name {
    description description;
    perfect-forward-secrecy {
      keys (group1 | group2 | group5 | group14 | group19 | group20);
    }
    proposals [ proposal-names ];
  }
}
rule rule-name {
  match-direction (input | output);
  term term-name {
    from {
      destination-address address;
      ipsec-inside-interface interface-name;
      source-address address;
    }
    then {
      anti-replay-window-size bits;
      backup-remote-gateway address;
    }
  }
}
```

```

clear-dont-fragment-bit;
dynamic {
    ike-policy policy-name;
    ipsec-policy policy-name;
}
dead-peer-detection {
    interval seconds ;
    threshold number ;
}
initiate-dead-peer-detection;
manual {
    direction (inbound | outbound | bidirectional) {
        authentication {
            algorithm (hmac-md5-96 | hmac-sha1-96 | hmac-sha-256-128);
            key (ascii-text key | hexadecimal key);
        }
        auxiliary-spi spi-value;
        encryption {
            algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
            key (ascii-text key | hexadecimal key);
        }
        protocol (ah | bundle | esp);
        spi spi-value;
    }
}
no-anti-replay;
remote-gateway address;
syslog;
tunnel-mtu bytes;
}
}
rule-set rule-set-name {
    [ rule rule-names ];
}
no-ipsec-tunnel-in-traceroute;
traceoptions {
    file {
        files number;
        size bytes;
    }
    flag flag;
    level level;
}

```

Related Documentation

- [Configuring Security Associations on page 385](#)
- [Configuring IKE Proposals on page 405](#)
- [Configuring IKE Policies on page 409](#)
- [Configuring IPsec Proposals on page 415](#)
- [Configuring IPsec Policies on page 420](#)
- [Configuring IPsec Rules on page 422](#)

- [Configuring IPsec Rule Sets on page 429](#)
- [Configuring Dynamic Endpoints for IPsec Tunnels on page 455](#)
- [Tracing Junos VPN Site Secure Operations on page 436](#)
- [Configuring Junos VPN Site Secure or IPSec VPN on page 507](#)

adaptive-services-pics

| | |
|---------------------------------|---|
| Syntax | <pre>adaptive-services-pics {
 traceoptions {
 file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i>> <size <i>size</i>> <world-readable
 no-world-readable>;
 flag <i>flag</i>;
 no-remote-trace;
 }
}</pre> |
| Hierarchy Level | [edit services] |
| Release Information | Statement introduced before Junos OS Release 7.4. The file option was added in Release 8.0. |
| Description | Define global services properties. |
| Options | The remaining statement is explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Tracing Services PIC Operations on page 48 |

address (Interfaces)

| | |
|---------------------------------|---|
| Syntax | <code>address address {
...
}</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the interface address. |
| Options | <i>address</i> —Address of the interface. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces. • Configuring the Logical Interface Address for the MLPPP Bundle on page 622 • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

address (Services NAT Pool)

| | |
|---------------------------------|--|
| Syntax | <code>address ip-prefix</prefix-length>;</code> |
| Hierarchy Level | [edit services nat pool <i>nat-pool-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
<i>prefix</i> option enhanced to support IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the NAT pool prefix value. |
| Options | <i>prefix</i> —Specify an IPv4 or IPv6 prefix value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Source and Destination Addresses Network Address Translation Overview on page 66 |

address-allocation

| | |
|--------------------------|--|
| Syntax | address-allocation round-robin; |
| Hierarchy Level | [edit services nat pool <i>pool-name</i>] |
| Release Information | Statement introduced in Junos OS Release 11.2. |
| Description | <p>When you use round-robin allocation, one port is allocated from each address in a range before repeating the process for each address in the next range. After ports have been allocated for all addresses in the last range, the allocation process wraps around and allocates the next unused port for addresses in the first range.</p> <p>Regardless of whether the round-robin method of allocation is addresses is enabled by using the address-allocation round-robin statement, round-robin allocation is enabled by default on MS-MICs and MS-MPCs.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Pools of Addresses and Ports for Network Address Translation Overview on page 67 |

address-range

| | |
|--------------------------|--|
| Syntax | address-range low <i>minimum-value</i> high <i>maximum-value</i> ; |
| Hierarchy Level | [edit services nat pool <i>nat-pool-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
<i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the NAT pool address range. |
| Options | <i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range.
<i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Source and Destination Addresses Network Address Translation Overview on page 66 |

aggregation

| | |
|---------------------------------|--|
| Syntax | <pre>aggregation {
 destination-prefix <i>prefix-value</i> destination-prefix-ipv6 <i>prefix-value</i>;
 source-prefix <i>prefix-value</i> source-prefix-ipv6 <i>prefix-value</i>;
}</pre> |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the type of data to be aggregated. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IDS Rules on page 355 |

allow-ip-options

Syntax `allow-ip-options [values];`

Hierarchy Level `[edit services stateful-firewall rule rule-name term term-name then]`

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure how the stateful firewall handles IP header information. This statement is optional.

Options *value*—Can be a set or range of numeric values, or one or more of the following predefined option types. You can enter either the option name or its numeric equivalent.

| Option Name | Numeric Value |
|---------------------|---------------|
| any | 0 |
| ip-security | 130 |
| ip-stream | 8 |
| loose-source-route | 3 |
| route-record | 7 |
| router-alert | 148 |
| strict-source-route | 9 |
| timestamp | 4 |

Usage Guidelines See “[Configuring Actions in Stateful Firewall Rules](#)” on page 334.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

allow-multicast

| | |
|---------------------------------|---|
| Syntax | allow-multicast; |
| Hierarchy Level | [edit services service-set <i>service-set-name</i>] |
| Release Information | Statement introduced in Junos OS Release 8.0. |
| Description | Allow multicast traffic to be sent to the Adaptive Services or Multiservices PIC. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Enabling Services PICs to Accept Multicast Traffic on page 38 |


allow-overlapping-nat-pools

| | |
|---------------------------------|---|
| Syntax | allow-overlapping-nat-pools; |
| Hierarchy Level | [edit services nat] |
| Release Information | Statement introduced with Junos OS Release 12.1. |
| Description | Specify that NAT source or destination pools can be shared between multiple service sets. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Sets for Network Address Translation on page 75 |

anti-replay-window-size (Services IPsec VPN)

| | |
|---------------------------------|--|
| Syntax | anti-replay-window-size <i>bits</i> ; |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced in Junos OS Release 10.0. |
| Description | Specify the size of the IPsec antireplay window. |
| Options | <i>bits</i> —Size of the antireplay window, in bits.
Default: 64 bits (AS PICs), 128 bits (Multiservices PICs and DPCs)
Range: 64 through 4096 bits |
| Required Privilege Level | admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Rules on page 422 |

anti-replay-window-size (Services Service Set)

| | |
|---------------------------------|--|
| Syntax | <code>anti-replay-window-size <i>bits</i>;</code> |
| Hierarchy Level | [edit services service-set <i>service-set-name</i> ipsec-vpn-options] |
| Release Information | Statement introduced in Junos OS Release 10.0. |
| Description | <p>Specify the size of the IPsec antireplay window. This statement is useful for dynamic endpoint tunnels for which you cannot configure the anti-replay-window-size statement at the [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] hierarchy level.</p> <p>For static IPsec tunnels, this statement sets the antireplay window size for all the static tunnels within this service set. If a particular tunnel needs a specific value for antireplay window size, set the anti-replay-window-size statement at the [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] hierarchy level. If antireplay check has to be disabled for a particular tunnel in this service set, set the no-anti-replay statement at the [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] hierarchy level.</p> |
| | <div>  <p>NOTE: The anti-replay-window-size and no-anti-replay settings at the [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] hierarchy level override the settings specified at the [edit services service-set <i>service-set-name</i> ipsec-vpn-options] hierarchy level.</p> </div> |
| Options | <p><i>bits</i>—Size of the antireplay window, in bits.</p> <p>Default: 64 bits (AS PICs), 128 bits (Multiservices PICs and DPCs)</p> <p>Range: 64 through 4096 bits</p> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IPsec Service Sets on page 430 • Configuring IPsec Rules |

app-mapping-timeout

| | |
|---------------------------------|---|
| Syntax | <code>app-mapping-timeout <i>app-mapping-timeout</i>;</code> |
| Hierarchy Level | [edit services nat pool <i>nat-pool-name</i>] |
| Release Information | <code>mapping-timeout</code> statement introduced in JUNOS Release 12.3. |
| Description | Specify the duration for address pooling paired (AP-P) mappings that use the specified NAT pool. If this option is not configured and a timeout value is configured for mapping-timeout , the timeout value configured for mapping-timeout is used. If neither option is specified, the default value of 300 seconds is used. |
| Options | <code>app-mapping-timeout</code> —Lifetime of AP-P mappings in seconds.
Default: 300
Range: 120 through 864,000 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Source and Destination Addresses Network Address Translation Overview on page 66 |

application

| | |
|---------------------------------|--|
| Syntax | <pre> application <i>application-name</i> { application-protocol <i>protocol-name</i>; destination-port <i>port-number</i>; icmp-code <i>value</i>; icmp-type <i>value</i>; inactivity-timeout <i>value</i>; protocol <i>type</i>; rpc-program-number <i>number</i>; snmp-command <i>command</i>; source-port <i>port-number</i>; ttl-threshold <i>number</i>; uuid <i>hex-value</i>; } </pre> |
| Hierarchy Level | [edit applications],
[edit applications application-set <i>application-set-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure properties of an application and whether to include it in an application set. |
| Options | <p><i>application-name</i>—Identifier of the application.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • ALG Descriptions on page 277 • Configuring Application Sets on page 303 • Configuring Application Protocol Properties on page 303 • Examples: Configuring Application Protocols on page 321 • Verifying the Output of ALG Sessions |

application-protocol

| | |
|----------------------------|--|
| Syntax | <code>application-protocol <i>protocol-name</i>;</code> |
| Hierarchy Level | [edit applications application <i>application-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
login options introduced in Junos OS Release 7.4.
ip option introduced in Junos OS Release 8.2. |
| Description | Identify the application protocol name. Application protocols are also called application layer gateways (ALGs). |
| Options | <p><i>protocol-name</i>—Name of the protocol. The following protocols are supported:</p> <ul style="list-style-type: none">bootp—Bootstrap protocoldce-rpc—DCE RPCdce-rpc-portmap—DCE RPC portmapdns—Domain Name Serviceexec—Remote Execution Protocolftp—File Transfer Protocolh323—H.323icmp—ICMPiiop—Internet Inter-ORB Protocolip—IPlogin—Loginnetbios—NetBIOSnetshow—NetShowpptp—Point-to-Point Tunneling Protocolrealaudio—RealAudiorpc—RPCrpc-portmap—RPC portmaprtsp—Real Time Streaming Protocolshell—Shellsip—Session Initiation Protocolsnmp—SNMPsqlnet—SQLNettalk—Talk Program |

tftp—Trivial File Transfer Protocol

traceroute—Traceroute

winframe—WinFrame

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [ALG Descriptions on page 277](#)
- [Configuring Application Sets on page 303](#)
- [Configuring Application Protocol Properties on page 303](#)
- [Examples: Configuring Application Protocols on page 321](#)
- [Verifying the Output of ALG Sessions](#)

application-profile

| | |
|--------------------------|--|
| Syntax | <pre>application-profile <i>profile-name</i> {
 ftp {
 data {
 dscp (<i>alias</i> <i>bits</i>);
 forwarding-class <i>class-name</i>;
 }
 }
 sip {
 video {
 dscp (<i>alias</i> <i>bits</i>);
 forwarding-class <i>class-name</i>;
 }
 voice {
 dscp (<i>alias</i> <i>bits</i>);
 forwarding-class <i>class-name</i>;
 }
 }
}</pre> |
| Hierarchy Level | [edit services cos],
[edit services cos rule <i>rule-name</i> term <i>term-name</i> then],
[edit services cos rule <i>rule-name</i> term <i>term-name</i> then (reflexive reverse)] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | Define or apply a CoS application profile. When you apply a CoS application profile in a CoS rule, terminate the profile name with a semicolon (;). |
| Options | <i>profile-name</i> —Identifier for the application profile.

The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">Configuring Application Profiles for Use as CoS Rule Actions on page 517 |

application-set

| | |
|---------------------------------|--|
| Syntax | <code>application-set <i>application-set-name</i> {
 application <i>application-name</i>;
}</code> |
| Hierarchy Level | [edit applications] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure one or more applications to include in an application set. |
| Options | <i>application-set-name</i> —Identifier of an application set. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • ALG Descriptions on page 277 • Configuring Application Sets on page 303 • Configuring Application Protocol Properties on page 303 • Examples: Configuring Application Protocols on page 321 • Verifying the Output of ALG Sessions |

application-sets (Services CoS)

| | |
|---------------------------------|---|
| Syntax | <code>applications-sets <i>set-name</i>;</code> |
| Hierarchy Level | [edit services cos rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | Define one or more target application sets. |
| Options | <i>set-name</i> —Name of the target application set. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Match Conditions In CoS Rules on page 515 |

application-sets (Services IDS)

| | |
|---------------------------------|---|
| Syntax | <code>application-sets set-name;</code> |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define one or more target application sets. |
| Options | <i>set-name</i> —Name of the target application set. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Match Conditions in IDS Rules on page 357 |

application-sets (Services NAT)

| | |
|---------------------------------|---|
| Syntax | <code>applications-sets set-name;</code> |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define one or more target application sets. |
| Options | <i>set-name</i> —Name of the target application set. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |

application-sets (Services Stateful Firewall)

| | |
|---------------------------------|---|
| Syntax | <code>applications-sets <i>set-name</i>;</code> |
| Hierarchy Level | [edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define one or more target application sets. |
| Options | <i>set-name</i> —Name of the target application set. |
| Usage Guidelines | See “ Configuring Match Conditions in Stateful Firewall Rules ” on page 332. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

applications (Services ALGs)

| | |
|---------------------------------|--|
| Syntax | <code>applications { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the applications used in services. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • ALG Descriptions on page 277 • Configuring Application Sets on page 303 • Configuring Application Protocol Properties on page 303 • Examples: Configuring Application Protocols on page 321 • Verifying the Output of ALG Sessions |

applications (Services CoS)

| | |
|---------------------------------|--|
| Syntax | <code>applications [<i>application-name</i>];</code> |
| Hierarchy Level | [edit services cos rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | Define one or more applications to which the CoS services apply. |
| Options | <i>application-name</i> —Name of the target application. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Match Conditions in a CoS Rule• Configuring Match Conditions In CoS Rules on page 515 |

applications (Services IDS)

| | |
|---------------------------------|---|
| Syntax | <code>applications [<i>application-names</i>];</code> |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define one or more applications to which IDS applies. |
| Options | <i>application-name</i> —Name of the target application. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Match Conditions in IDS Rules on page 357 |


applications (Services NAT)

| | |
|---------------------------------|---|
| Syntax | <code>applications [<i>application-names</i>];</code> |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define one or more application protocols to which the NAT services apply. |
| Options | <i>application-name</i> —Name of the target application. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Network Address Translation Rules Overview on page 69 |

applications (Services Stateful Firewall)

| | |
|---------------------------------|---|
| Syntax | <code>applications [<i>application-names</i>];</code> |
| Hierarchy Level | [edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define one or more applications to which the stateful firewall services apply. |
| Options | <i>application-name</i> —Name of the target application. |
| Usage Guidelines | See “ Configuring Match Conditions in Stateful Firewall Rules ” on page 332. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

authentication (Services IPsec VPN)

| | |
|---------------------------------|---|
| Syntax | <pre>authentication {
 algorithm (hmac-md5-96 hmac-sha1-96 hmac-sha-256-128);
 key (ascii-text key hexadecimal key);
}</pre> |
| Hierarchy Level | [edit services ipsec-vpn rule rule-name term term-name then manual direction direction] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure IPsec authentication parameters for a manual security association (SA). |
| Options | <p>algorithm—Hash algorithm that authenticates packet data. The algorithm can be one of the following:</p> <ul style="list-style-type: none">• hmac-md5-96—Produces a 128-bit digest.• hmac-sha1-96—Produces a 160-bit digest.• hmac-sha-256-128—Produces a 256-bit digest, truncated to 128 bits. <hr/> <div> NOTE: hmac-sha-256-128 is not supported on MS-MIC and MS-MPC.</div> <hr/> <p>key—Type of authentication key. The key can be one of the following:</p> <ul style="list-style-type: none">• ascii-text key—ASCII text key. For hmac-md5-96, the key is 16 ASCII characters; for hmac-sha1-96, the key is 20 ASCII characters.• hexadecimal key—Hexadecimal key. For hmac-md5-96, the key is 32 hexadecimal characters; for hmac-sha1-96, the key is 40 hexadecimal characters. |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Security Associations on page 385 |

authentication-algorithm (Services IKE)

| | |
|---------------------------------|---|
| Syntax | authentication-algorithm (md5 sha1 sha-256); |
| Hierarchy Level | [edit services ipsec-vpn ike proposal <i>proposal-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
sha-256 option added in Junos OS Release 7.6. |
| Description | Configure the Internet Key Exchange (IKE) hash algorithm that authenticates packet data. |
| Options | md5 —Produces a 128-bit digest.
sha1 —Produces a 160-bit digest.
sha-256 —Produces a 256-bit digest.
sha-384 —Produces a 384-bit digest. |
| Required Privilege Level | admin —To view this statement in the configuration.
admin-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IKE Proposals on page 405 |

authentication-algorithm (Services IPsec)

| | |
|----------------------------|--|
| Syntax | authentication-algorithm (hmac-md5-96 hmac-sha-256-128 hmac-sha1-96); |
| Hierarchy Level | [edit services ipsec-vpn ipsec proposal <i>ipsec-proposal-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the IPsec hash algorithm that authenticates packet data. |



NOTE: Keep the following points in mind when you configure the authentication algorithm in an IPsec proposal:

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, an error occurs and the tunnel is not established in this scenario. For example, if one end of the tunnel contains router 1 configured with the authentication algorithm as hmac-sha-256-128 and the other end of the tunnel contains router 2 configured with the authentication algorithm as hmac-md5-96, the VPN tunnel is not established.
- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, and when one end of the tunnel contains two IPsec proposals to check whether a less secure algorithm is selected or not, an error occurs and the tunnel is not established. For example, if you configure two authentication algorithms for an IPsec proposal as hmac-sha-256-128 and hmac-md5-96 on one end of the tunnel, router 1, and if you configure the algorithm for an IPsec proposal as hmac-md5-96 on the other end of the tunnel, router 2, the tunnel is not established and the number of proposals mismatch.
- When you configure two IPsec proposals at both ends of a tunnel, such as the authentication-algorithm hmac-sha-256-128 and authentication-algorithm hmac-md5-96 statements at the [edit services ipsec-vpn ipsec proposal *proposal-name*] hierarchy level on one of the tunnel, router 1 (with the algorithms in two successive statements to specify the order), and the authentication-algorithm hmac-md5-96 and authentication-algorithm hmac-sha-256-128 statements at the [edit services ipsec-vpn ipsec proposal *proposal-name*] hierarchy level on one of the tunnel, router 2 (with the algorithms in two successive statements to specify the order, which is the reverse order of router 1), the tunnel is established in this combination as expected because the number of proposals is the same on both ends and they contain the same set of algorithms. However, the authentication algorithm selected is hmac-md5-96 and not the stronger algorithm of hmac-sha-256-128. This method of selection of the algorithm occurs because the first matching proposal is selected. Also, for a default proposal, regardless of whether the router supports the Advanced Encryption Standard (AES) encryption algorithm, the 3des-cbc algorithm is chosen and not the aes-cfb algorithm, which is because of the first algorithm in

the default proposal being selected. In the sample scenario described here, on router 2, if you reverse the order of the algorithm configuration in the proposal so that it is the same order as the one specified on router 1, hmac-sha-256-128 is selected as the authentication method.

- You must be aware of the order of proposals in an IPsec policy at the time of configuration if you want the matching of proposals to happen in a certain order of preference, such as the strongest algorithm to be considered first when a match is made when both policies from the two peers have a proposal.

| | |
|---------------------------------|--|
| Options | <p>hmac-md5-96—Produces a 128-bit digest.</p> <p>hmac-sha-256-128—Produces a 256-bit digest.</p> <p>hmac-sha1-96—Produces a 160-bit digest.</p> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IPsec Proposals on page 415 |

authentication-method (Services IPsec VPN)

| | |
|---------------------------------|--|
| Syntax | authentication-method (pre-shared-keys rsa-signatures); |
| Hierarchy Level | [edit services ipsec-vpn ike proposal <i>proposal-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure an IKE authentication method. |
| Options | <p>rsa-signatures—Public key algorithm (supports encryption and digital signatures).</p> <p>pre-shared-keys—A key derived from an out-of-band mechanism; the key authenticates the exchange.</p> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IKE Proposals on page 405 |

auxiliary-spi (Services IPsec VPN)

| | |
|---------------------------------|---|
| Syntax | <code>auxiliary-spi <i>spi-value</i>;</code> |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then manual direction <i>direction</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure an auxiliary Security Parameter Index (SPI) for a manual SA. Use the auxiliary SPI when you configure the protocol statement to use the bundle option. |
| Options | <i>spi-value</i> —An arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet).
Range: 256 through 16,639 |
| Required Privilege Level | admin —To view this statement in the configuration.
admin-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Security Associations on page 385 |

backup-remote-gateway

| | |
|---------------------------------|--|
| Syntax | <code>backup-remote-gateway <i>address</i>;</code> |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the backup remote address to which the IPsec traffic is directed when the primary remote gateway is down. Configuring this statement also enables the dead peer detection (DPD) protocol. |
| Options | <i>address</i> —Backup remote IPv4 or IPv6 address. |
| Required Privilege Level | admin —To view this statement in the configuration.
admin-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Rules on page 422 |

bundle

| | |
|---------------------------------|---|
| Syntax | <code>bundle (lsq-fpc/pic/port ...);</code> |
| Hierarchy Level | [edit interfaces lsq-fpc/pic/port unit logical-unit-number family mlppp] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Associate the voice services interface with the logical interface it is joining. |
| Options | lsq-fpc/pic/port —Name of the voice services interface you are linking. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Voice Services Bundles with MLPPP Encapsulation on page 626 |

by-destination

| | |
|---------------------------------|--|
| Syntax | <pre>by-destination { hold-time <i>seconds</i>; maximum <i>number</i>; packets <i>number</i>; rate <i>number</i>; }</pre> |
| Hierarchy Level | [edit services ids rule rule-name term term-name then session-limit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Apply limit to sessions based on numbers generated from the configured destination (IP or subnet) or application. |
| Options | <p>hold-time <i>seconds</i>—Length of time for which to stop all new flows once the rate of events exceeds the threshold set by one or more of the maximum, packets, or rate statements.</p> <p>maximum <i>number</i>—Maximum number of open sessions per application or IP address.</p> <p>packets <i>number</i>—Maximum peak packets per second per application or IP address.</p> <p>rate <i>number</i>—Maximum number of sessions per second per application or IP address.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Actions in IDS Rules on page 358 |

by-pair

| | |
|--------------------------|--|
| Syntax | <pre>by-pair {
 hold-time <i>seconds</i>;
 maximum <i>number</i>;
 packets <i>number</i>;
 rate <i>number</i>;
}</pre> |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> then session-limit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Apply limit to paired stateful firewall and NAT flows (forward and reverse). |
| Options | <p>hold-time <i>seconds</i>—Length of time for which to stop all new flows once the rate of events exceeds the threshold set by one or more of the maximum, packets, or rate statements.</p> <p>maximum <i>number</i>—Maximum number of open sessions per application or IP address.</p> <p>packets <i>number</i>—Maximum peak packets per second per application or IP address.</p> <p>rate <i>number</i>—Maximum number of sessions per second per application or IP address.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Actions in IDS Rules on page 358 |

by-source

| | |
|---------------------------------|--|
| Syntax | <pre>by-source { hold-time <i>seconds</i>; maximum <i>number</i>; packets <i>number</i>; rate <i>number</i>; }</pre> |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> then session-limit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Apply limit to sessions based on numbers generated from the configured source (IP or subnet) or application. |
| Options | <p>hold-time <i>seconds</i>—Length of time for which to stop all new flows once the rate of events exceeds the threshold set by one or more of the maximum, packets, or rate statements.</p> <p>maximum <i>number</i>—Maximum number of open sessions per application or IP address.</p> <p>packets <i>number</i>—Maximum peak packets per second per application or IP address.</p> <p>rate <i>number</i>—Maximum number of sessions per second per application or IP address.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> Configuring Actions in IDS Rules on page 358 |

bypass-traffic-on-exceeding-flow-limits

| | |
|---------------------------------|---|
| Syntax | bypass-traffic-on-exceeding-flow-limits; |
| Hierarchy Level | [edit services service-set <i>service-set-name</i> service-set-options] |
| Release Information | Statement introduced in Junos OS Release 10.1. |
| Description | Enable packets to bypass without creating a new session when the flow in the service set exceeds the limit that is set by the max-flows statement at the [edit services service-set <i>service-set-name</i>] hierarchy level. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> Configuring Service Sets to be Applied to Services Interfaces on page 31 |

bypass-traffic-on-pic-failure

| | |
|---------------------------------|--|
| Syntax | bypass-traffic-on-pic-failure; |
| Hierarchy Level | [edit services service-set <i>service-set-name</i> service-set-options] |
| Release Information | Statement introduced in Junos OS Release 10.1. |
| Description | <p>When the MultiServices PIC configured for a service set is either administratively taken offline or undergoes a failure, all the traffic entering the configured interface with an IDP service set would be dropped without notification. To avoid this traffic loss, include the bypass-traffic-on-pic-failure statement. When this statement is configured, the affected packets are forwarded in the event of a MultiServices PIC failure or offlining, as though interface-style services were not configured.</p> <p>This issue applies only to Junos Application Aware (previously known as Dynamic Application Awareness) configurations with IDP service sets.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Sets to be Applied to Services Interfaces on page 31 |

cgn-pic

| | |
|---------------------------------|--|
| Syntax | cgn-pic; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options] |
| Release Information | Statement introduced in Junos OS Release 11.2. |
| Description | Restrict usage of the service PIC to carrier-grade NAT (CGN) or associated services (intrusion detection, stateful firewall, and software). All memory is available for CGN or related services and can be used for CGN scaling. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 61 |

cisco-interoperability

| | |
|---------------------------------|--|
| Syntax | cisco-interoperability send-lip-remove-link-for-link-reject; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | FRF.16 interoperability settings. |
| Options | send-lip-remove-link-for-link-reject —Send Link Integrity Protocol remove link when an add-link rejection message is received. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring SONET APS Interoperability with Cisco Systems FRF.16 on page 547 |

class

| | |
|---------------------------------|--|
| Syntax | <pre>class { alg-logs; ids-logs; nat-logs; packet-logs; pcp-logs; session-logs <open close>; stateful-firewall-logs ; }</pre> |
| Hierarchy Level | [edit services service-set <i>service-set-name</i> syslog host <i>hostname</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2. |
| Description | Set the class of applications to be logged to the system log. |
| Options | <p><i>class-name</i>—Enter one of the following values:</p> <ul style="list-style-type: none">• alg-logs—Log application-level gateway events.• ids-logs—Log intrusion detection system events.• nat-logs—Log Network Address Translation events.• packet-logs—Log general packet-related events.• session-logs—Log session open and close events.• session-logs open—Log session open events only.• session-logs close—Log session close events. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• See Configuring System Logging for Service Sets on page 47. |

clear-dont-fragment-bit (Interfaces GRE Tunnels)

| | |
|---------------------------------|--|
| Syntax | clear-dont-fragment-bit; |
| Hierarchy Level | [edit interfaces gr-fpc/pic/port unit logical-unit-number],
[edit logical-systems logical-system-name interfaces gr-fpc/pic/port unit logical-unit-number] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>Clear the Don't Fragment (DF) bit on all IP version 4 (IPv4) packets entering the generic routing encapsulation (GRE) tunnel on Adaptive Services (AS) or Multiservices interfaces. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. The statement is supported only on MX Series routers and all M Series routers except the M320 router.</p> <p>When you configure the clear-dont-fragment-bit statement on an interface with the MPLS protocol family enabled, you must specify an MTU value. This MTU value must not be greater than maximum supported value, which is 9192.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Enabling Fragmentation on GRE Tunnels on page 1208 |

clear-dont-fragment-bit (Services IPsec VPN)

| | |
|---------------------------------|---|
| Syntax | clear-dont-fragment-bit; |
| Hierarchy Level | [edit services ipsec-vpn rule rule-name term term-name then] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Clear the do not fragment (DF) bit on all IP version 4 (IPv4) packets entering the IPsec tunnel. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IPsec Rules on page 422 |

clear-dont-fragment-bit (Services Service Set)

| | |
|---------------------------------|--|
| Syntax | clear-dont-fragment-bit; |
| Hierarchy Level | [edit services service-set <i>service-set-name</i> ipsec-vpn-options] |
| Release Information | Statement introduced in Junos OS Release 10.0. |
| Description | <p>Clear the Don't Fragment (DF) bit on all IP version 4 (IPv4) packets entering the IPsec tunnel. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. This statement is useful for dynamic endpoint tunnels, for which you cannot configure the clear-dont-fragment-bit statement at the [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] hierarchy level.</p> <p>For static IPsec tunnels, setting this statement clears the DF bit on packets entering all the static tunnels within this service set. If you want to clear the DF bit on packets entering a specific tunnel, set the clear-dont-fragment-bit statement at the [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] hierarchy level.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Service Sets on page 430• Configuring IPsec Rules |

clear-ike-sas-on-pic-restart

| | |
|---------------------------------|---|
| Syntax | clear-ike-sas-on-pic-restart; |
| Hierarchy Level | [edit services ipsec-vpn] |
| Release Information | Statement introduced in Junos OS Release 8.5. |
| Description | Clear IKE security associations (SAs) when the corresponding PIC restarts or is taken offline. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Security Associations on page 385 |

clear-ipsec-sas-on-pic-restart

| | |
|---------------------------------|---|
| Syntax | clear-ipsec-sas-on-pic-restart; |
| Hierarchy Level | [edit services ipsec-vpn] |
| Release Information | Statement introduced in Junos OS Release 9.2. |
| Description | Clear IPsec security associations (SAs) when the corresponding PIC restarts or is taken offline. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Security Associations on page 385 |

compression

| | |
|---------------------------------|--|
| Syntax | <pre> compression { rtp { f-max-period <i>number</i>; maximum-contexts <i>number</i> <force>; port { minimum <i>port-number</i>; maximum <i>port-number</i>; } queues [<i>queue-numbers</i>]; } } </pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>Configure the compression properties for voice services traffic.</p> <p>The remaining statements are described separately.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Compression of Voice Traffic on page 623 |

compression-device (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>compression-device <i>interface-name</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced in Junos OS Release 7.5. |
| Description | Specify the compression interface for voice services traffic. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Compression Interface with PPP Encapsulation on page 626 |

copy-dont-fragment-bit (Services IPsec VPN)

| | |
|---------------------------------|---|
| Syntax | <code>copy-dont-fragment-bit;</code> |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced in Junos OS Release 14.1. |
| Description | Copy the do not fragment (DF) bit value to only the outer header and not modify the inner header of the IPsec packet. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. This functionality is supported on MX Series routers with MS-MICs and MS-MPCs. These settings apply for static endpoint tunnels and not for dynamic tunnels, for which you need to include the copy-dont-fragment-bit statement at the [edit services service-set <i>service-set-name</i> ipsec-vpn-options] hierarchy level to copy the DF bit value to only the outer header of the packet in a static IPsec tunnel. This functionality is supported on MX Series routers with MS-MICs and MS-MPCs. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Rules on page 422 |

copy-dont-fragment-bit (Services Set)

| | |
|---------------------------------|--|
| Syntax | <code>copy-dont-fragment-bit;</code> |
| Hierarchy Level | [edit services service-set <i>service-set-name</i> ipsec-vpn-options] |
| Release Information | Statement introduced in Junos OS Release 14.1. |
| Description | Copy the do not fragment (DF) bit value to only the outer header and not modify the inner header of the IPsec packet in dynamic endpoint tunnels. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. This functionality is supported on MX Series routers with MS-MICs and MS-MPCs. These settings apply for dynamic endpoint tunnels and not for static tunnels, for which you need to include the copy-dont-fragment-bit statement at the [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] hierarchy level to copy the DF bit value to only the outer header of the packet in a static IPsec tunnel. This functionality is supported on MX Series routers with MS-MICs and MS-MPCs. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IPsec Service Sets on page 430 • Configuring IPsec Rules |

data (FTP)

| | |
|---------------------------------|--|
| Syntax | <pre>data { dscp (<i>alias</i> <i>bits</i>); forwarding-class <i>class-name</i>; }</pre> |
| Hierarchy Level | [edit services cos application-profile <i>profile-name</i> ftp] |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | Set the appropriate dscp and forwarding-class value for FTP data. |
| Default | By default, the system will not alter the DSCP or forwarding class for FTP data traffic. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Application Profiles • video (Application Profile) on page 1532 • voice (Application Profile) on page 1533 |

dead-peer-detection (Services IPsec VPN)

| | |
|---------------------------------|--|
| Syntax | <code>dead-peer-detection {
 interval <i>seconds</i>;
 threshold <i>number</i>;
}</code> |
| Hierarchy Level | [edit services ipsec-vpn <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Sets dead peer detection options when dead peer detection has been enabled with the initiate-dead-peer-detection command. The dead-peer-detection options are used for IKEv1 security associations (SAs) but not for IKEv2 SAs. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | security—To view this statement in the configuration.
security-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Rules on page 422 |

description (Services IPsec VPN)

| | |
|---------------------------------|---|
| Syntax | <code>description <i>description</i>;</code> |
| Hierarchy Level | [edit services ipsec-vpn ike policy <i>policy-name</i>],
[edit services ipsec-vpn ike proposal <i>proposal-name</i>],
[edit services ipsec-vpn ipsec policy <i>policy-name</i>],
[edit services ipsec-vpn ipsec proposal <i>proposal-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the text description for an IKE or IPsec policy or proposal. |
| Required Privilege Level | admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• description on page 1334• Configuring IPsec Proposals on page 415• Configuring IPsec Policies on page 420 |

destination-address (Services CoS)

| | |
|---------------------------------|---|
| Syntax | <code>destination-address (address any-unicast) <except>;</code> |
| Hierarchy Level | [edit services cos rule rule-name term term-name from] |
| Release Information | Statement introduced in Junos OS Release 8.1.
address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the destination address for rule matching. |
| Options | address —Destination IPv4 or IPv6 address or prefix value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Match Conditions in a CoS Rule • Configuring Match Conditions In CoS Rules on page 515 |

destination-address (Services IDS)

| | |
|---------------------------------|--|
| Syntax | <code>destination-address (address any-unicast) <except>;</code> |
| Hierarchy Level | [edit services ids rule rule-name term term-name from] |
| Release Information | Statement introduced before Junos OS Release 7.4.
address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the destination address for rule matching. |
| Options | address —Destination IPv4 or IPv6 address or prefix value.

any-unicast —Any unicast packet.

except —(Optional) Exempt the specified address, prefix, or unicast packets from rule matching. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Match Conditions in IDS Rules on page 357 |

destination-address (Services IPsec VPN)

| | |
|---------------------------------|---|
| Syntax | <code>destination-address <i>address</i>;</code> |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the destination address for rule matching. |
| Options | <i>address</i> —Destination IP address. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Rules on page 422 |

destination-address (Services NAT)

| | |
|---------------------------------|--|
| Syntax | <code>destination-address (<i>address</i> any-unicast) <except>;</code> |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4.
<i>any-unicast</i> and <i>except</i> options introduced in Junos OS Release 7.6.
<i>address</i> option enhanced to support IPv6 and addresses in Junos OS Release 8.5. |
| Description | Specify the destination address for rule matching. |
| Options | <i>address</i> —Destination IPv4 or IPv6 address or prefix value.
<i>any-unicast</i> —Any unicast packet.
<i>except</i> —(Optional) Prevent the specified address, prefix, or unicast packets from being translated. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |

destination-address (Services Stateful Firewall)

| | |
|---------------------------------|---|
| Syntax | <code>destination-address (address any-unicast) <except>;</code> |
| Hierarchy Level | [edit services stateful-firewall rule rule-name term term-name from] |
| Release Information | Statement introduced before Junos OS Release 7.4.
any-unicast and except options introduced in Junos OS Release 7.6.
address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the destination address for rule matching. |
| Options | address —Destination IPv4 or IPv6 address or prefix value. Using a value of 0::0/0 with IPv6 is not allowed for M-Series and MX-Series routers.

any-unicast —Match all unicast packets.

except —(Optional) Exclude the specified address, prefix, or unicast packets from rule matching. |
| Usage Guidelines | See “ Configuring Match Conditions in Stateful Firewall Rules ” on page 332. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

destination-address-range (Services IDS)

| | |
|---------------------------------|--|
| Syntax | <code>destination-address-range low minimum-value high maximum-value <except>;</code> |
| Hierarchy Level | [edit services ids rule rule-name term term-name from] |
| Release Information | Statement introduced in Junos OS Release 7.6.
minimum-value and maximum-value options enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the destination address range for rule matching. |
| Options | minimum-value —Lower boundary for the IPv4 or IPv6 address range.

maximum-value —Upper boundary for the IPv4 or IPv6 address range.

except —(Optional) Exempt the specified address range from rule matching. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Match Conditions in IDS Rules on page 357 |

destination-address-range (Services NAT)

| | |
|--------------------------|--|
| Syntax | destination-address-range low <i>minimum-value</i> high <i>maximum-value</i> <except>; |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 7.6.
<i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the destination address range for rule matching.

If the translation-type statement in the then statement of the nat rule is set to stateful-nat-64 , the destination address range for rule matching must be within the range specified by the destination-prefix statement in the then statement. |
| Options | <i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range.

<i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range.

except —(Optional) Prevent the specified address range from being translated. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |

destination-address-range (Services Stateful Firewall)

| | |
|--------------------------|---|
| Syntax | destination-address-range low <i>minimum-value</i> high <i>maximum-value</i> <except>; |
| Hierarchy Level | [edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 7.6.
<i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the destination address range for rule matching. |
| Options | <i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range.

<i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range.

except —(Optional) Exclude the specified address range from rule matching. |
| Usage Guidelines | See “ Configuring Match Conditions in Stateful Firewall Rules ” on page 332. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

destination-pool

| | |
|---------------------------------|--|
| Syntax | <code>destination-pool <i>nat-pool-name</i>;</code> |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the destination address pool for translated traffic. |
| Options | <i>nat-pool-name</i> —Destination pool name. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |

destination-port

| | |
|----------------------------|---|
| Syntax | <code>destination-port <i>port-value</i>;</code> |
| Hierarchy Level | [edit applications application application-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number. |
| Options | <i>port-value</i> —Identifier for the port or range of ports. For a complete list of supported application destination port requirements, see “Configuring Source and Destination Ports” on page 309 .
Range: 1 through 65,535 |



NOTE: If you specify a value of 0 as a destination port or beginning of a destination report range, you will receive the following error:

```
'application application-name'  
  TCP Destination Port 0 Invalid  
error: configuration check-out failed
```

| | |
|---------------------------------|---|
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• ALG Descriptions on page 277• Configuring Application Sets on page 303• Configuring Application Protocol Properties on page 303• Examples: Configuring Application Protocols on page 321• Verifying the Output of ALG Sessions• Two-Way Active Measurement Protocol Overview |

destination-port range

| | |
|---------------------------------|--|
| Syntax | destination-port range <i>high</i> <i>low</i> ; |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Specify the destination port range for rule matching. |
| Options | <i>high</i> —Upper limit of port range for matching.
<i>low</i> —Lower limit of port range for matching. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Port Forwarding for Static Destination Address Translation on page 183 |

destination-prefix (Services IDS)

| | |
|---------------------------------|---|
| Syntax | destination-prefix <i>prefix-value</i> ; |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> then aggregation] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the prefix value for destination IPv4 address aggregation. |
| Options | <i>prefix-value</i> —Integer value.
Range: 1 through 32 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Actions in IDS Rules on page 358 |

destination-prefix (Services NAT)

| | |
|--------------------------|---|
| Syntax | destination-prefix <i>destination-prefix</i> ; |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated] |
| Release Information | Statement introduced in Junos OS Release 7.6.
<i>destination-prefix</i> option enhanced to support IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the destination prefix for translated traffic. |
| Options | <i>destination-prefix</i> —IPv4 or IPv6 destination prefix value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |

destination-prefix-ipv6

| | |
|--------------------------|---|
| Syntax | destination-prefix-ipv6 <i>prefix</i> ; |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> then aggregation] |
| Release Information | Statement introduced in Junos OS Release 8.5. |
| Description | Specify the prefix value for destination IPv6 address aggregation. |
| Options | <i>prefix-value</i> —Integer value.
Range: 1 through 128 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Actions in IDS Rules on page 358 |

destination-prefix-list (Services CoS)

| | |
|---------------------------------|--|
| Syntax | <code>destination-prefix-list <i>list-name</i> <except>;</code> |
| Hierarchy Level | [edit services cos rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 8.2. |
| Description | Specify the destination prefix list for rule matching. You configure the prefix list by including the prefix-list statement at the [edit policy-options] hierarchy level. |
| Options | <p><i>list-name</i>—Destination prefix list.</p> <p>except—(Optional) Exclude the specified prefix list from rule matching.</p> |
| Usage Guidelines | See “ Configuring Match Conditions In CoS Rules ” on page 515. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> |

destination-prefix-list (Services IDS)

| | |
|---------------------------------|--|
| Syntax | <code>destination-prefix-list <i>list-name</i> <except>;</code> |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 8.2. |
| Description | Specify the destination prefix list for rule matching. You configure the prefix list by including the prefix-list statement at the [edit policy-options] hierarchy level. |
| Options | <p><i>list-name</i>—Destination prefix list.</p> <p>except—(Optional) Exclude the specified prefix list from rule matching.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> Configuring Match Conditions in IDS Rules on page 357 |

destination-prefix-list (Services NAT)

| | |
|---------------------------------|---|
| Syntax | <code>destination-prefix-list <i>list-name</i> <except>;</code> |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 8.2. |
| Description | <p>Specify the destination prefix list for rule matching. You configure the prefix list by including the prefix-list statement at the [edit policy-options] hierarchy level.</p> <p>If the translation-type statement in the then statement of the nat rule is set to stateful-nat-64, the destination prefix list for rule matching must be within the range specified by the destination-prefix statement in the then statement.</p> |
| Options | <p><i>list-name</i>—Destination prefix list.</p> <p>except—(Optional) Exclude the specified prefix list from rule matching.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> |


destination-prefix-list (Services Stateful Firewall)

| | |
|---------------------------------|--|
| Syntax | <code>destination-prefix-list <i>list-name</i> <except>;</code> |
| Hierarchy Level | [edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 8.2. |
| Description | <p>Specify the destination prefix list for rule matching. You configure the prefix list by including the prefix-list statement at the [edit policy-options] hierarchy level.</p> |
| Options | <p><i>list-name</i>—Destination prefix list.</p> <p>except—(Optional) Exclude the specified prefix list from rule matching.</p> |
| Usage Guidelines | See “ Configuring Match Conditions in Stateful Firewall Rules ” on page 332. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> |

destined-port

| | |
|---------------------------------|---|
| Syntax | <code>destined-port <i>port id</i>;</code> |
| Hierarchy Level | [edit services nat port-forwarding <i>map-name</i>] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Specify the port from where traffic has to be forwarded. |
| Options | <i>port id</i> —The destination port number from where traffic will be forwarded. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• port-forwarding on page 1433• translated-port on page 1518 |

deterministic-port-block-allocation

| | |
|--|--|
| Syntax | <pre>deterministic-port-block-allocation {
 block-size <i>block-size</i>;
 include-boundary-addresses;
}</pre> |
| Hierarchy Level | [edit services nat pool <i>pool-name</i> port] |
| Release Information | Statement introduced in Junos OS Release 12.1. |
| Description | Configure algorithm-based allocation of blocks of destination ports. By specifying this method, you ensure that an incoming (source) IP address and port always map to the same destination IP address and port block, thus eliminating the need for logging address translations. |
| Options | <i>block-size</i> —Maximum number of ports that can be allocated to a user. |
| <hr/> | |
| <div> NOTE: When a block-size of 0 is specified, block size is calculated according to the formula: $(64512 * \text{Number of IP addresses in the NAT Pool}) / \text{Number of subscribers}$ where</div> <ul style="list-style-type: none">• 64512 is derived from (65535 - 1023) because the regular port assignments start from 1024.• Number of subscribers is derived from the from clause of the applicable NAT rule. <hr/> | |
| Default: 256 | |
| Range: 0 through 32,000 | |
| include-boundary-addresses —(Optional) Specifies that the lowest and highest addresses in the source address range of a NAT rule should be translated when the NAT pool is used. | |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Deterministic Port Block Allocation on page 180 |

dh-group

| | |
|---------------------------------|---|
| Syntax | dh-group (group1 group2 group5 group14 group19 group20); |
| Hierarchy Level | [edit services ipsec-vpn ike proposal <i>proposal-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the IKE Diffie-Hellman prime modulus group to use for performing the new Diffie-Hellman exchange. |
| Options | <p>group1—768-bit.</p> <p>group2—1024-bit.</p> <p>group5—1536-bit.</p> <p>group14—2048-bit.</p> <p>group19—256-bit random Elliptic Curve Group.</p> <p>group20—384-bit random Elliptic Curve Group.</p> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IKE Proposals on page 405 |

dial-options

| | |
|---------------------------------|---|
| Syntax | <pre>dial-options {
 ipsec-interface-id <i>name</i>;
 l2tp-interface-id <i>name</i>;
 (shared dedicated);
}</pre> |
| Hierarchy Level | <pre>[edit interfaces sp-fpc/pic/port unit <i>logical-unit-number</i>],
[edit interfaces si-fpc/pic/port unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces sp-fpc/pic/port unit
 <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces si-fpc/pic/port unit <i>logical-unit-number</i>]</pre> |
| Release Information | Statement introduced before Junos OS Release 7.4.
The [edit ...si-...] hierarchy levels introduced in Junos OS Release 11.4. |
| Description | Specify the options for configuring logical interfaces for group and user sessions in L2TP or IPsec dynamic endpoint tunneling. |
| Options | <p>dedicated—(LNS on M Series routers and MX Series routers only) Specify that a logical interface can host only one session at a time.</p> <p>ipsec-interface-id <i>name</i>—(M Series routers only) Interface identifier for group of dynamic peers. This identifier must be replicated at the [edit access profile <i>name</i> client * ike] hierarchy level.</p> <p>l2tp-interface-id <i>name</i>—Interface identifier that must be replicated at the [edit access profile <i>name</i>] hierarchy level.</p> <p>shared—(LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 646• <i>Configuring Dynamic Endpoints for IPsec Tunnels</i>• <i>Configuring Options for the LNS Inline Services Logical Interface</i> |

direction

| | |
|---------------------------------|--|
| Syntax | <pre> direction (inbound outbound bidirectional) { protocol (ah bundle esp); spi spi-value; auxiliary-spi spi-value; authentication { algorithm (hmac-md5-96 hmac-sha1-96); key (ascii-text key hexadecimal key); } encryption { algorithm algorithm; key (ascii-text key hexadecimal key); } } </pre> |
| Hierarchy Level | [edit services ipsec-vpn rule rule-name term term-name then manual] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the direction in which manual SAs are applied. |
| Options | <p>bidirectional—Apply the SA in both directions.</p> <p>inbound—Apply the SA on inbound traffic.</p> <p>outbound—Apply the SA on outbound traffic.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IPsec Rules on page 422 |

dns-alg-pool

| | |
|---------------------------------|--|
| Syntax | dns-alg-pool <i>dns-alg-pool</i> ; |
| Hierarchy Level | [edit services nat rule rule-name term term-name then translated] |
| Release Information | Statement introduced in Junos OS Release 10.4. |
| Description | Specify the Network Address Translation (NAT) pool for destination translation. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Network Address Translation Rules Overview on page 69 |

dns-alg-prefix

| | |
|---------------------------------|--|
| Syntax | <code>dns-alg-prefix <i>dns-alg-prefix</i>;</code> |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated] |
| Release Information | Statement introduced in Junos OS Release 10.4. |
| Description | Set the Domain Name System (DNS) application-level gateway (ALG) 96-bit prefix for mapping IPv4 addresses to IPv6 addresses. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |

drop-member-traffic (Aggregated Multiservices)

| | |
|---------------------------------|---|
| Syntax | <code>drop-member-traffic {
 rejoin-timeout <i>rejoin-timeout</i>;
}</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> load-balancing-options member-failure-options] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | <p>Specify whether the broadband gateway should drop traffic to a Multiservices PIC when it fails.</p> <p>For many-to-one (N:1) high availability (HA) for service applications like Network Address Translation (NAT), this configuration is valid only when two or more Multiservices PICs have failed.</p> <p>The remaining statement is explained separately.</p> |
| Default | If this statement is not configured, then the default behavior is to drop member traffic with a rejoin timeout of 120 seconds. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• member-failure-options (Aggregated Multiservices) on page 1408• Understanding Aggregated Multiservices Interfaces on page 599• Example: Configuring an Aggregated Multiservices Interface (AMS) on page 608 |

ds-lite

| | |
|---------------------------------|---|
| Syntax | <pre>ds-lite <i>ds-lite-softwire-concentrator</i> { auto-update-mtu; copy-dscp; flow-limit <i>flow-limit</i> session-limit-per-prefix <i>session-limit-per-prefix</i>; mtu-v6 <i>mtu-v6</i>; softwire-address <i>softwire-address</i>; }</pre> |
| Hierarchy Level | [edit services softwire softwire-concentrator] |
| Release Information | <p>Statement introduced in Junos OS Release 10.4.</p> <p>auto-update-mtu option introduced in Junos OS Release 10.4.</p> <p>copy-dscp option introduced in Junos OS Release 11.2.</p> <p>mtu-v6 option introduced in Junos OS Release 10.4.</p> <p>softwire-address option introduced in Junos OS Release 10.4.</p> |
| Description | Configure settings for a DS-Lite concentrator used to process IPv4 packets encapsulated in IPv6. |
| Options | <p><i>ds-lite-softwire-concentrator</i>—Name applied to a DS-Lite softwire concentrator.</p> <p>auto-update-mtu—This option is not currently supported.</p> <p>copy-dscp—Copy DSCP information to IPv4 headers during decapsulation.</p> <p><i>flow-limit</i>—Maximum number of IPv4 flows per softwire.
 Range: 0 through 16384 flows</p> <p><i>mtu-v6</i>—Maximum transmission unit (MTU), in bytes, for encapsulating IPv4 packets into IPv6. If the final length is greater than the configured value, the IPv6 packet is fragmented.
 Range: 0 through 9192 bytes</p> <p><i>session-limit-per-prefix</i>—Maximum number of sessions per B4 subnet prefix. (0 through 16384).
 Range: 0 through 16384 sessions</p> <p><i>softwire-address</i>—Address of the DS-Lite softwire concentrator.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring a DS-Lite Softwire Concentrator on page 227 |

dscp

| | |
|---------------------------------|--|
| Syntax | <code>dscp (<i>alias</i> <i>bits</i>);</code> |
| Hierarchy Level | [edit services cos application-profile <i>profile-name</i> ftp data],
[edit services cos application-profile <i>profile-name</i> sip (video voice)],
[edit services cos rule <i>rule-name</i> term <i>term-name</i> then],
[edit services cos rule <i>rule-name</i> term <i>term-name</i> then (reflexive reverse)] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | Define the Differentiated Services code point (DSCP) mapping that is applied to the packets. |
| Options | <i>alias</i> —Name assigned to a set of CoS markers.

<i>bits</i> —Mapping value in the packet header. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Actions in CoS Rules on page 516. |

dynamic

| | |
|---------------------------------|--|
| Syntax | <pre>dynamic {
 ike-policy <i>policy-name</i>;
 ipsec-policy <i>policy-name</i>;
}</pre> |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define a dynamic IPsec SA. |
| Options | <i>ike-policy policy-name</i> —Name of the IKE policy. This statement is optional for the non-preshared-key authentication method. For digital signature-based authentication, this statement is optional and the default policy is used if none is supplied.

<i>ipsec-policy policy-name</i> —Name of the IPsec policy. This statement is optional and the default policy is used if none is supplied. |
| Required Privilege Level | admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Security Associations on page 385 |

ecmp-alb

Syntax `ecmp-alb {
 apply-groups;
 apply-groups-except;
 tolerance;
 }`

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 14.2.

Description Enable adaptive load balancing for equal-cost multipath (ECMP) next hops.



NOTE: The `ecmp-alb` statement can be enabled only when the `[edit chassis network-services enhanced-ip]` statement is configured.

Options **apply-groups**—Specify the groups from which to inherit configuration data.

apply-groups-except—Specify the groups from which configuration data should not be inherited.

tolerance—Specify the adaptive tolerance in percentage.

Default: 20%.

Required Privilege interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

ei-mapping-timeout

| | |
|---------------------------------|--|
| Syntax | mapping-timeout <i>seconds</i> ; |
| Hierarchy Level | [edit services nat pool <i>nat-pool-name</i>] |
| Release Information | ei-mapping-timeout statement introduced in JUNOS Releases 12.3. |
| Description | Specify the duration for endpoint independent translations that use the specified NAT pool. This includes endpoint-independent mapping (EIM) and endpoint-independent filtering (EIF). |
| Options | seconds —Lifetime of endpoint independent mappings in seconds.
Default: 300
Range: 120 through 864,000 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Configuration Overview on page 65 |

eif-flow-limit

| | |
|---------------------------------|--|
| Syntax | eif-flow-limit <i>number-of-flows</i> |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated secure-nat-mapping] |
| Release Information | Statement introduced in Junos OS Release 12.3 |
| Description | Specify the maximum number of inbound flows allowed on EIF mapping to the configured value. This limit is per EIF mapping and is per given instance of time. For example, if eif-flow-limit is configured as n, then only n inbound connections are allowed at a given instance of time. The n+1 and subsequent connections arriving when n connections are alive are dropped . A new inbound connection is allowed only when one of the n connections times out or is closed. This limit is applied for all type of flows. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Protecting CGN Devices Against Denial of Service (DOS) Attacks on page 241 |

enable-rejoin (aggregated Multiservices)

| | |
|---------------------------------|--|
| Syntax | enable-rejoin; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> load-balancing-options member-failure-options redistribute-all-traffic] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | <p>Enable the failed member to rejoin the aggregated Multiservices (AMS) interface after the member comes back online.</p> <p>For many-to-one (N:1) high availability (HA) for service applications like Network Address Translation (NAT), this configuration allows the failed members to rejoin the pool of active members automatically.</p> |
| Default | If you do not configure this option, then the failed members do not automatically rejoin the ams interface even after coming back online. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • redistribute-all-traffic (Aggregated Multiservices) on page 1442 • Understanding Aggregated Multiservices Interfaces on page 599 • Example: Configuring an Aggregated Multiservices Interface (AMS) on page 608 |

encapsulation

| | |
|---------------------------------|--|
| Syntax | <code>encapsulation type;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the logical link-layer encapsulation type. |
| Options | atm-mlppp-llc —For ATM2 IQ physical interfaces only, use Multilink Point-to-Point Protocol (MLPPP) over AAL5 LLC encapsulation.

frame-relay-ppp —For Frame Relay circuits, use Frame Relay PPP encapsulation.

multilink-ppp —By default, voice services logical interfaces use MLPPP encapsulation. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Encapsulation for Voice Services on page 625• <i>Junos OS Network Interfaces Library for Routing Devices</i> |

encryption

| | |
|----------------------------|---|
| Syntax | <pre> encryption { algorithm <i>algorithm</i>; key (ascii-text <i>key</i> hexadecimal <i>key</i>); } </pre> |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then manual direction <i>direction</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
aes-128-cbc , aes-192-cbc , and aes-256-cbc options added in Junos OS Release 7.6. |
| Description | Configure an encryption algorithm and key for manual SA. |

Options **algorithm**—Type of encryption algorithm. The algorithm can be one of the following:

- **des-cbc**—Has a block size of 8 bytes (64 bits); the key size is 48 bits long.
- **3des-cbc**—Has a block size of 8 bytes (64 bits); the key size is 192 bits long.
- **aes-128-cbc**—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- **aes-192-cbc**—Advanced Encryption Standard (AES) 192-bit encryption algorithm.
- **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption algorithm.



NOTE: For **3des-cbc**, the first 8 bytes should differ from the second 8 bytes, and the second 8 bytes should be the same as the third 8 bytes.

key—Type of encryption key. The key can be one of the following:

- **ascii-text**—ASCII text key. Following are the key lengths, in ASCII characters, for the different encryption options:
 - **des-cbc** option, 8 ASCII characters
 - **3des-cbc** option, 24 ASCII characters
 - **aes-128-cbc** option, 16 ASCII characters
 - **aes-192-cbc** option, 24 ASCII characters
 - **aes-256-cbc** option, 32 ASCII characters
- **hexadecimal**—Hexadecimal key. Following are the key lengths, in hexadecimal characters, for the different encryption options:
 - **des-cbc** option, 16 hexadecimal characters
 - **3des-cbc** option, 48 hexadecimal characters
 - **aes-128-cbc** option, 32 hexadecimal characters
 - **aes-192-cbc** option, 48 hexadecimal characters

- **aes-256-cbc** option, 64 hexadecimal characters

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- [Configuring Security Associations on page 385](#)

encryption-algorithm (Services IPsec VPN)

Syntax encryption-algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);

Hierarchy Level [edit [services](#) ipsec-vpn [ike proposal](#) *proposal-name*],
[edit [services](#) ipsec-vpn [ipsec proposal](#) *proposal-name*]

Release Information Statement introduced before Junos OS Release 7.4.
aes-128-cbc, **aes-192-cbc**, and **aes-256-cbc** options added in Junos OS Release 7.6.

Description Configure an IKE or IPsec encryption algorithm.

Options **3des-cbc**—Has a block size of 24 bytes; the key size is 192 bits long.

des-cbc—Has a block size of 8 bytes; the key size is 48 bits long.

aes-128-cbc—Advanced Encryption Standard (AES) 128-bit encryption algorithm.

aes-192-cbc—Advanced Encryption Standard (AES) 192-bit encryption algorithm.

aes-256-cbc—Advanced Encryption Standard (AES) 256-bit encryption algorithm.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring Security Associations on page 385](#)

establish-tunnels

| | |
|---------------------------------|--|
| Syntax | establish-tunnels (immediately on-traffic); |
| Hierarchy Level | [edit services ipsec-vpn] |
| Release Information | Statement introduced in Release 8.5 of Junos OS. |
| Description | Specify when IKE is activated: immediately after VPN information is configured and configuration changes are committed, or only when data traffic flows. In the second case, IKE needs to be negotiated with the peer gateway. |
| Options | <ul style="list-style-type: none"> • immediately—IKE is activated immediately after VPN configuration and configuration changes are committed. • on-traffic—IKE is activated only when data traffic flows and must to be negotiated. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • IPsec Hierarchy Level on page 1297 |

f-max-period

| | |
|---------------------------------|--|
| Syntax | f-max-period <i>number</i> ; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression <i>rtp</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression <i>rtp</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the maximum number of compressed packets allowed between the transmission of full headers in a compressed Real-time Transport Protocol (RTP) traffic stream. |
| Options | <i>number</i> —Maximum number of packets.
Default: 256 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Compression of Voice Traffic on page 623 |

facility-override (Service Sets)

| | |
|---------------------------------|--|
| Syntax | <code>facility-override <i>facility-name</i>;</code> |
| Hierarchy Level | [edit services service-set service-set-name syslog host hostname] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Override the default facility for system log reporting. |
| Options | <p><i>facility-name</i>—Name of the facility that overrides the default assignment. Valid entries are:</p> <ul style="list-style-type: none"><code>authorization</code><code>daemon</code><code>ftp</code><code>kernel</code><code>local0</code> through <code>local7</code><code>user</code> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring System Logging for Service Sets on page 47 |

facility-override (System Log Reporting)

| | |
|---------------------------------|--|
| Syntax | <code>facility-override <i>facility-name</i>;</code> |
| Hierarchy Level | [edit services l2tp tunnel-group <i>group-name</i> syslog host <i>hostname</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Override the default facility for system log reporting. |
| Options | <p><i>facility-name</i>—Name of the facility that overrides the default assignment. Valid entries include:</p> <ul style="list-style-type: none"> authorization daemon ftp kernel local0 through local7 user |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring System Logging of L2TP Tunnel Activity on page 644 |

family (Aggregated Multiservices)

| | |
|---------------------------------|---|
| Syntax | <code>family <i>family</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>interface-unit-number</i>] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Configure protocol family information for the logical interface. |
| Options | <i>family</i> —Protocol family. Currently, only one option, inet (IP version 4 suite), is supported. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • unit (Aggregated Multiservices) on page 1527 • Understanding Aggregated Multiservices Interfaces on page 599 • Example: Configuring an Aggregated Multiservices Interface (AMS) on page 608 |

family (Interfaces)

| | |
|--------------------------|--|
| Syntax | <pre>family inet {
 address address {
 ...
 }
 service {
 input {
 [service-set service-set-name <service-filter filter-name>];
 post-service-filter filter-name;
 }
 output {
 [service-set service-set-name <service-filter filter-name>];
 }
 }
}</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure protocol family information for the logical interface. |
| Options | family —Protocol family. Valid settings for service interfaces include inet (IPv4) and mpls .

The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces.• Configuring the Address and Domain for Services Interfaces on page 45• <i>Junos OS Network Interfaces Library for Routing Devices</i> |

family (Voice Services)

| | |
|---------------------------------|--|
| Syntax | <pre>family (inet mlppp ...) { address address { ... } bundle interface-name; }</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure protocol family information for the logical interface. |
| Options | <p><i>family</i>—Protocol family:</p> <ul style="list-style-type: none"> • <i>inet</i>—IP version 4 • <i>mlppp</i>—MLPPP <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces. • Configuring Network Interfaces for Voice Services on page 626 • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

force-entry

| | |
|---------------------------------|---|
| Syntax | (force-entry ignore-entry); |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify handling of entries in the IDS events cache: <ul style="list-style-type: none">• force-entry—Ensure that the entry has a permanent place in the IDS cache after one event is registered.• ignore-entry—Ensure that all IDS events are ignored. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Actions in IDS Rules on page 358 |

forwarding-class (Services CoS)

| | |
|---------------------------------|--|
| Syntax | forwarding-class <i>class-name</i> ; |
| Hierarchy Level | [edit services cos application-profile <i>profile-name</i> ftp data],
[edit services cos application-profile <i>profile-name</i> sip (video voice)],
[edit services cos rule <i>rule-name</i> term <i>term-name</i> then],
[edit services cos rule <i>rule-name</i> term <i>term-name</i> then (reflexive reverse)] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | Define the forwarding class to which packets are assigned. |
| Options | <i>class-name</i> —Name of the target application. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Actions in CoS Rules on page 516. |

forwarding-class (Services CoS Fragmentation Properties)

| | |
|---------------------------------|--|
| Syntax | <pre>forwarding-class <i>class-name</i> { (<i>fragment-threshold bytes</i> <i>no-fragmentation</i>); <i>multilink-class number</i>; }</pre> |
| Hierarchy Level | [edit class-of-service fragmentation-maps] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>For link services IQ (lsq) interfaces only, define a forwarding class name and associated fragmentation properties within a fragmentation map.</p> <p>The fragment-threshold and no-fragmentation statements are mutually exclusive.</p> |
| Default | If you do not include this statement, the traffic in forwarding class <i>class-name</i> is fragmented. |
| Options | <p><i>class-name</i>—Name of the forwarding class.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 527 |

fragment-threshold (Forwarding Class Maps)

| | |
|---------------------------------|---|
| Syntax | <code>fragment-threshold <i>bytes</i>;</code> |
| Hierarchy Level | [edit class-of-service fragmentation-maps forwarding-class <i>class-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services IQ (lsq) interfaces only, set the fragmentation threshold for an individual forwarding class. |
| Default | If you do not include this statement, the fragmentation threshold you set at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] or [edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options] hierarchy level is the default for all forwarding classes. If you do not set a maximum fragment size anywhere in the configuration, packets are fragmented if they exceed the smallest maximum transmission unit (MTU) of all the links in the bundle. |
| Options | bytes —Maximum size, in bytes, for multilink packet fragments. Any nonzero value must be a multiple of 64 bytes.
Range: 128 through 16,320 bytes |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 527 |

fragment-threshold (Interfaces LSQ)

| | |
|---------------------------------|--|
| Syntax | <code>fragment-threshold <i>bytes</i>;</code> |
| Hierarchy Level | [edit interfaces <i>lsq-fpc/pic/port</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>lsq-fpc/pic/port</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For voice services interfaces, set the fragmentation threshold, in bytes. |
| Options | bytes —Maximum size, in bytes, for multilink packet fragments. The value must be a multiple of 64 bytes, because zero is also a multiple of 64 bytes.
Range: 128 through 16,320 bytes
Default: 0 bytes (no fragmentation) |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Delay-Sensitive Packet Interleaving on page 624 |

fragmentation-map

| | |
|---------------------------------|---|
| Syntax | <code>fragmentation-map <i>map-name</i>;</code> |
| Hierarchy Level | [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services IQ (lsq) interfaces only, associate a fragmentation map with a multilink PPP interface or MLFR FRF.16 DLCI. |
| Default | If you do not include this statement, traffic in all forwarding classes is fragmented. |
| Options | map-name —Name of the fragmentation map. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 527 |

fragmentation-maps

| | |
|---------------------------------|---|
| Syntax | <pre>fragmentation-maps {
 map-name {
 forwarding-class class-name {
 (fragment-threshold bytes no-fragmentation);
 multilink-class number;
 }
 }
}</pre> |
| Hierarchy Level | [edit class-of-service] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services IQ (lsq) interfaces only, define fragmentation properties for individual forwarding classes. |
| Default | If you do not include this statement, traffic in all forwarding classes is fragmented. |
| Options | <p><i>map-name</i>—Name of the fragmentation map.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 527 |

from (Services CoS)

| | |
|---------------------------------|---|
| Syntax | <pre> from { application-sets set-name; applications [application-names]; destination-address address; destination-prefix-list list-name <except>; source-address address; source-prefix-list list-name <except>; } </pre> |
| Hierarchy Level | [edit services cos rule rule-name term term-name] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | Specify input conditions for a CoS term. |
| Options | <p>For information on match conditions, see the description of firewall filter match conditions in the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring CoS Rules on page 514 |

from (Services IDS)

| | |
|---------------------------------|--|
| Syntax | <pre>from {
 application-sets set-name;
 applications [application-names];
 destination-address (address any-unicast) <except>;
 destination-address-range low minimum-value high maximum-value <except>;
 source-address (address any-unicast) <except>;
 source-address-range low minimum-value high maximum-value <except>;
}</pre> |
| Hierarchy Level | [edit services ids rule rule-name term term-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify input conditions for the IDS term. |
| Options | <p>For information on match conditions, see the description of firewall filter match conditions in the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Match Conditions in IDS Rules on page 357 |

from (Services IPsec VPN)

| | |
|---------------------------------|--|
| Syntax | <pre> from { destination-address address; ipsec-inside-interface interface-name; source-address address; } </pre> |
| Hierarchy Level | [edit services ipsec-vpn rule rule-name term term-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify input conditions for the IPsec term. |
| Options | <p>For information on match conditions, see the description of firewall filter match conditions in the Junos OS Routing Policy Configuration Guide..</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IPsec Rules on page 422 |

from (Services HCM)

| | |
|---------------------------------|--|
| Syntax | <pre> from { url-list url-list-name; url url_identifier { host hostname; request-url page-name; } } </pre> |
| Hierarchy Level | [edit services hcm url-rule url-rule-name term term-num] |
| Release Information | Statement introduced in Junos OS Release 12.1. |
| Description | Specify input conditions for the HCM term. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |

from (Services NAT)

| | |
|---------------------------------|--|
| Syntax | <pre>from {
 application-sets set-name;
 applications [application-names];
 destination-address (address any-unicast) <except>;
 destination-address-range low minimum-value high maximum-value <except>;
 source-address address (address any-unicast) <except>;
 source-address-range low minimum-value high maximum-value <except>;
}</pre> |
| Hierarchy Level | [edit services nat rule rule-name term term-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify input conditions for the NAT term. |
| Options | <p>For information on match conditions, see the description of firewall filter match conditions in the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |

from (Services Stateful Firewall)

Syntax from {
 [application-sets](#) *set-name*;
 [applications](#) [*application-names*];
 [destination-address](#) (*address* | any-unicast) <except>;
 [destination-address-range](#) low *minimum-value* high *maximum-value* <except>;
 [destination-prefix-list](#) *list-name* <except>;
 [source-address](#) (*address* | any-unicast) <except>;
 [source-address-range](#) low *minimum-value* high *maximum-value* <except>;
 [source-prefix-list](#) *list-name* <except>;
 }

Hierarchy Level [edit [services](#) stateful-firewall [rule](#) *rule-name* [term](#) *term-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify input conditions for a stateful firewall term.

Options For information on match conditions, see the description of firewall filter match conditions in the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

The remaining statements are explained separately.

Usage Guidelines See “[Configuring Stateful Firewall Rules](#)” on page 331.


Required Privilege interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.


ftp (Services CoS)

| | |
|--------------------------|--|
| Syntax | <pre>ftp {
 data {
 dscp (alias bits);
 forwarding-class class-name;
 }
}</pre> |
| Hierarchy Level | [edit services cos application-profile <i>profile-name</i> ftp] |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | Set the appropriate dscp and forwarding-class value for FTP. |
| Default | By default, the system does not alter the DSCP or forwarding class for FTP traffic. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Application Profiles• sip (Application Profile) |

hello-interval

| | |
|---|---|
| Syntax | hello-interval <i>seconds</i> ; |
| Hierarchy Level | [edit services l2tp tunnel-group <i>name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the keepalive timer for L2TP tunnels. |
| <div> NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.</div> | |
| Options | seconds —Interval, in seconds, after which the server sends a hello message if no messages are received. A value of 0 means that no hello messages are sent.
Default: 60 seconds |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Timers for L2TP Tunnels on page 643• Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces |

hide-avps

| | |
|--|---|
| Syntax | hide-avps; |
| Hierarchy Level | [edit services l2tp tunnel-group name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Hide L2TP attribute-value pairs if the secret shared between the two ends of the tunnel is known. |
| <div>  NOTE: This statement is not supported for L2TP LNS on MX Series routers. </div> | |
| Default | Attribute-value pairs that can be hidden are exposed, even if the secret information is known. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Hiding Attribute-Value Pairs for L2TP Tunnels on page 644 |

high-availability-options (aggregated Multiservices)

Syntax high-availability-options {
 many-to-one {
 preferred-backup *preferred-backup*;
 }
 }

Hierarchy Level [edit interfaces *interface-name* load-balancing-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the high availability options for the aggregated Multiservices (AMS) interface. For service applications, if only the load-balancing feature is being used, then this configuration is optional.

For many-to-one (N:1) high availability support for service applications like Network Address Translation (NAT), the preferred backup Multiservices PIC, in hot standby mode, backs up one or more (N) active Multiservices PICs.



NOTE: In both cases, if one of the active Multiservices PICs goes down, then the backup replaces it as the active Multiservices PIC. When the failed PIC comes back up, it becomes the new backup. This is called floating backup.

The remaining statements are explained separately.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation

- [load-balancing-options on page 1393](#)
- [Understanding Aggregated Multiservices Interfaces on page 599](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 608](#)

host (L2TP)

| | |
|---------------------------------|---|
| Syntax | <pre>host <i>hostname</i> {
 <i>services severity-level</i>;
 <i>facility-override facility-name</i>;
 <i>log-prefix prefix-value</i>;
}</pre> |
| Hierarchy Level | [edit services l2tp tunnel-group <i>group-name</i> syslog] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the hostname for the system logging utility. |
| Options | <p>hostname—Name of the system logging utility host machine. This can be the local Routing Engine or an external server address.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring System Logging of L2TP Tunnel Activity on page 644 |

host (service-set)

Syntax `host hostname {
 class {
 alg-logs;
 ids-logs;
 nat-logs;
 packet-logs;
 pcp-logs;
 session-logs <open | close>;
 stateful-firewall-logs ;
 }
 facility-override facility-name;
 interface-service prefix-value;
 log-prefix prefix-value
 port port-number
 services severity-level;
 source-address source-address
 }`

Hierarchy Level [edit [services service-set service-set-name syslog](#)]

Release Information Statement introduced before Junos OS Release 7.4.
 class option introduced in Junos OS Release 13.2.

Description Specify the hostname for the system logging utility.

Options *hostname*—Name of the system logging utility host machine.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Configuring System Logging for Service Sets on page 47](#)

host (Services HCM)

| | |
|----------------------------|---|
| Syntax | <code>host <i>hostname</i>;</code> |
| Hierarchy Level | [edit services hcm url-rule <i>url-rule-name</i> term <i>term-num</i> from url <i>url_identifier</i>] |
| Release Information | Statement introduced in Junos OS Release 12.1. |
| Description | Specify a hostname for the matching URL for the term . A match for that term is considered when a URL matches the hostname and the request-URL within the same term. |
| Options | hostname —Name of the host for the URL rule. |
| Required Privilege | interface—To view this statement in the configuration. |
| Level | interface-control—To add this statement to the configuration. |

hot-standby

| | |
|------------------------------|--|
| Syntax | <code>hot-standby;</code> |
| Hierarchy Level | [edit interfaces <i>rlsnumber</i> redundancy-options],
[edit interfaces <i>rlsnumber:number</i> redundancy-options]
[edit interfaces <i>rspnumber</i> redundancy-options]
[edit interfaces <i>rmsnumber</i> redundancy-options] |
| Release Information | Statement introduced in Junos OS Release 7.6. |
| Description | For one-to-one AS, rsp, or rms redundancy configurations, specify that the failure detection and recovery must take place in less than 5 seconds. For FRF.15 (MLFR) and FRF.16 (MFR) configuration, specify the switch over time of 5 seconds and less for FRF.15 and a maximum of 10 seconds for FRF.16. |
| Required Privilege | interface—To view this statement in the configuration. |
| Level | interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 548 • Configuring AS or Multiservices PIC Redundancy on page 41 |

icmp-code

| | |
|---------------------------------|---|
| Syntax | <code>icmp-code value;</code> |
| Hierarchy Level | [edit applications application application-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Internet Control Message Protocol (ICMP) code value. |
| Options | value —The ICMP code value. For a complete list, see “ Configuring the ICMP Code and Type ” on page 307. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• ALG Descriptions on page 277• Configuring Application Sets on page 303• Configuring the ICMP Code and Type on page 307• Examples: Configuring Application Protocols on page 321• Verifying the Output of ALG Sessions |

icmp-type

| | |
|---------------------------------|---|
| Syntax | <code>icmp-type value;</code> |
| Hierarchy Level | [edit applications application application-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | ICMP packet type value. |
| Options | value —The ICMP type value, such as echo or echo-reply . For a complete list, see “ Configuring the ICMP Code and Type ” on page 307. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• ALG Descriptions on page 277• Configuring Application Sets on page 303• Configuring the ICMP Code and Type on page 307• Examples: Configuring Application Protocols on page 321• Verifying the Output of ALG Sessions |

ids-rules

| | |
|---------------------------------|---|
| Syntax | (ids-rules <i>rule-name</i> ids-rule-sets <i>rule-set-name</i>); |
| Hierarchy Level | [edit services service-set service-set-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the intrusion detection service (IDS) rules or rule set included in this service set. You can configure multiple rules, but only one rule set for each service. |
| Options | <i>rule-name</i> —Identifier for the collection of terms that constitute this rule.
<i>rule-set-name</i> —Identifier for the set of rules to be included. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Rules on page 36 |

ignore-entry

See [force-entry](#)

ike

Syntax `ike {`
 `proposal proposal-name {`
 `authentication-algorithm (md5 | sha1 | sha-256);`
 `authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);`
 `description description;`
 `dh-group (group1 | group2 | group5 | group14);`
 `encryption-algorithm algorithm;`
 `lifetime-seconds seconds;`
 `}`
 `policy policy-name {`
 `description description;`
 `local-certificate identifier;`
 `local-id (ipv4_addr ipv4-address | ipv6_addr ipv6-address | key-id identifier);`
 `version (1 | 2);`
 `mode (aggressive | main);`
 `pre-shared-key (ascii-text key | hexadecimal key);`
 `proposals [proposal-names];`
 `remote-id {`
 `any-remote-id;`
 `ipv4_addr [values];`
 `ipv6_addr [values];`
 `key_id [values];`
 `}`
 `}`
 `}`

Hierarchy Level [edit [services](#) ipsec-vpn]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure IKE.

The statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation

- [Configuring IKE Proposals on page 405](#)
- [Configuring IKE Policies on page 409](#)

ike-access-profile

| | |
|---------------------------------|--|
| Syntax | <code>ike-access-profile <i>profile-name</i>;</code> |
| Hierarchy Level | [edit services service-set service-set-name ipsec-vpn-options] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Define the access profile for the IPsec traffic on dynamic tunnels. |
| Options | <i>profile-name</i> —Identifier for access profile, which must match the name configured at the [edit <code>access profile name client * ike</code>] hierarchy level. |
| Required Privilege Level | admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring Dynamic Endpoints for IPsec Tunnels</i> • Configuring IPsec Service Sets on page 430 |

inactivity-timeout

| | |
|---------------------------------|---|
| Syntax | <code>inactivity-timeout <i>seconds</i>;</code> |
| Hierarchy Level | [edit applications application application-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Inactivity timeout period, in seconds. |
| Options | <i>seconds</i> —Length of time the application is inactive before it times out.
Default: 30 seconds |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • ALG Descriptions on page 277 • Configuring Application Sets on page 303 • Configuring the Inactivity Timeout Period on page 312 • Examples: Configuring Application Protocols on page 321 • <i>Verifying the Output of ALG Sessions</i> |

initiate-dead-peer-detection

| | |
|---------------------------------|---|
| Syntax | <code>initiate-dead-peer-detection;</code> |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced in Junos OS Release 9.2. |
| Description | Enable triggering of dead peer detection (DPD) hello messages to the remote peer for the specified tunnel. |
| Required Privilege Level | system—To view this statement in the configuration.
system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Rules on page 422• dead-peer-detection on page 1334• backup-remote-gateway on page 1322• Configuring Destination Addresses for Dead Peer Detection |

input (Interfaces)

| | |
|---------------------------------|---|
| Syntax | <pre>input {
 service-set <i>service-set-name</i> <service-filter <i>filter-name</i>>;
 post-service-filter <i>filter-name</i>;
}</pre> |
| Hierarchy Level | [edit interface <i>interface-name</i> unit <i>logical-unit-number</i> family inet service],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the input service sets and filters to be applied to traffic. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Applying Filters and Services to Interfaces on page 38 |

interface-service

| | |
|---------------------------------|--|
| Syntax | <pre>interface-service {
 service-interface <i>name</i>;
}</pre> |
| Hierarchy Level | [edit services service-set <i>service-set-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the device name for the interface service Physical Interface Card (PIC). |
| Options | service-interface <i>name</i> —Name of the service device associated with the interface-wide service set. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Sets to be Applied to Services Interfaces on page 31 |

interfaces (Aggregated Multiservices)

```
Syntax  interfaces interface-name {
        load-balancing-options {
            high-availability-options {
                many-to-one {
                    preferred-backup preferred-backup;
                }
            }
        }
        member-failure-options {
            drop-member-traffic {
                rejoin-timeout rejoin-timeout;
            }
            redistribute-all-traffic {
                enable-rejoin;
            }
        }
        member-interface interface-name;
    }
    unit interface-unit-number {
        family family;
    }
}
```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the aggregated Multiservices (AMS) interface. The AMS interface provides the infrastructure for load balancing and high availability (HA).



NOTE: The interfaces must be valid aggregated Multiservices interfaces (*ams*)—for example, *ams0* or *ams1*, and so on. The *ams* infrastructure is supported only in chassis with Trio-based modules and Multiservices Dense Port Concentrators (MS-DPCs).

The remaining statements are explained separately.

Options *interface-name*—Name of the aggregated Multiservices interface (*ams*)—for example, *ams0* or *ams1*, and so on.

Required Privilege Level *interface*—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Load Balancing on AMS Infrastructure on page 605](#)
- [Understanding Aggregated Multiservices Interfaces on page 599](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 608](#)

interfaces (Voice Services)

| | |
|---------------------------------|---|
| Syntax | <code>interfaces { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure interfaces on the router. |
| Default | The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

interval

| | |
|---------------------------------|--|
| Syntax | <code>interval seconds;</code> |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then dead-peer-detection] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Specify the amount of time that the peer waits for traffic from its destination peer before sending a dead-peer-detection (DPD) request packet. (The interval value is used for IKEv1 security associations (SAs) but not for IKEv2 SAs.) |
| Options | seconds —Number of seconds that the peer waits before sending a DPD request packet.
Range: 1 through 180 seconds
Default: 2 seconds |
| Required Privilege Level | security—To view this statement in the configuration.
security-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IPsec Rules on page 422 |

ipsec (Services IPsec VPN)

Syntax ipsec {
 proposal *proposal-name* {
 authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
 description *description*;
 encryption-algorithm *algorithm*;
 lifetime-seconds *seconds*;
 protocol (ah | esp | bundle);
 }
 policy *policy-name* {
 description *description*;
 perfect-forward-secrecy {
 keys (group1 | group2);
 }
 proposals [*proposal-names*];
 }
 }

Hierarchy Level [edit [services](#) ipsec-vpn]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure IPsec.

The statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation • [Configuring Security Associations on page 385](#)

ipsec-inside-interface

Syntax ipsec-inside-interface *interface-name*;

Hierarchy Level [edit [services](#) ipsec-vpn **rule** *rule-name* **term** *term-name* **from**]

Release Information Statement introduced in Junos OS Release 7.4.

Description Specify the interface name for next-hop-style service sets. This value is also implicitly generated in dynamic endpoint tunneling.

Options *interface-name*—Service interface for internal network.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Configuring IPsec Rules on page 422](#)
 • [Configuring Dynamic Endpoints for IPsec Tunnels on page 455](#)

ipsec-vpn-options

| | |
|---------------------------------|---|
| Syntax | <pre>ipsec-vpn-options { anti-replay-window-size <i>bits</i>; clear-dont-fragment-bit; ike-access-profile <i>profile-name</i>; local-gateway <i>address</i>; no-anti-replay; passive-mode-tunneling; trusted-ca [<i>ca-profile-names</i>]; tunnel-mtu <i>bytes</i>; }</pre> |
| Hierarchy Level | [edit services service-set service-set-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify IP Security (IPsec) service options. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Service Rules on page 36 |

ipsec-vpn-rules

| | |
|---------------------------------|---|
| Syntax | (ipsec-vpn-rules rule-name ipsec-vpn-rule-sets rule-set-name); |
| Hierarchy Level | [edit services service-set service-set-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the IPsec rules or rule set included in this service set. You can configure multiple rules, but only one rule set for each service. |
| Options | <p>rule-name—Identifier for the collection of terms that constitute this rule.</p> <p>rule-set-name—Identifier for the set of rules to be included.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Service Rules on page 36 |

ipv6-multicast-interfaces

| | |
|---------------------------------|---|
| Syntax | ipv6-multicast-interfaces (all <i>interface-name</i>) {
disable;
} |
| Hierarchy Level | [edit services nat],
[edit services software] |
| Release Information | Statement introduced in Junos OS Release 9.1. |
| Description | Enable multicast filters on Ethernet interfaces when IPv6 NAT is used for neighbor discovery. |
| Options | all —Enable filters on all interfaces.

disable —Disable filters on the specified interfaces.

<i>interface-name</i> —Enable filters on a specific interface only. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPv6 Multicast Interfaces |

l2tp-access-profile

| | |
|---------------------------------|--|
| Syntax | l2tp-access-profile <i>profile-name</i> ; |
| Hierarchy Level | [edit services l2tp tunnel-group name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the profile used to validate all L2TP connection requests to the local gateway address. |
| Options | <i>profile-name</i> —Identifier for the L2TP connection profile. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Access Profiles for L2TP Tunnel Groups on page 642• Configuring an L2TP Access Profile on the LNS |

learn-sip-register

| | |
|---------------------------------|--|
| Syntax | learn-sip-register; |
| Hierarchy Level | [edit applications application application-name] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Activate SIP register to accept potential incoming SIP calls. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • ALG Descriptions on page 277 • Configuring Application Sets on page 303 • Configuring SIP on page 303 • Examples: Configuring Application Protocols on page 321 • Verifying the Output of ALG Sessions |

lifetime-seconds (Services IPsec VPN)

| | |
|---------------------------------|---|
| Syntax | lifetime-seconds <i>seconds</i> ; |
| Hierarchy Level | [edit services ipsec-vpn ike proposal proposal-name],
[edit services ipsec-vpn ipsec proposal proposal-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the lifetime of an IKE or IPsec SA. This statement is optional. |
| Options | <i>seconds</i> —Lifetime
Default: 3600 seconds (IKE); 28,800 seconds (IPsec)
Range: 180 through 86,400 |
| Required Privilege Level | system—To view this statement in the configuration.
system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Security Associations on page 385 |

link-layer-overhead

| | |
|---------------------------------|---|
| Syntax | link-layer-overhead <i>percent</i> ; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options],
[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services IQ (lsq) interfaces only, configure the percentage of total bundle bandwidth to be set aside for link-layer overhead. Link-layer overhead accounts for the bit stuffing on serial links. Bit stuffing is used to prevent data from being interpreted as control information. Overhead resulting from link-layer encapsulation and framing is computed automatically. |
| Options | percent —Percentage of total bundle bandwidth to be set aside for link-layer overhead.
Range: 0 through 50 percent
Default: 0 percent |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring CoS Scheduling Queues on Logical LSQ Interfaces on page 523 |

load-balance

| | |
|---------------------------------|--|
| Syntax | load-balance {
per-packet;
random;
} |
| Hierarchy Level | [edit policy-options policy-statement <i>policy-name</i> then] |
| Release Information | Statement introduced in Junos OS Release 14.2. |
| Description | Specify the type of load balancing of an equal-cost multipath (ECMP) in the forwarding table. |
| Options | per-packet —Load-balance on a per-packet basis.

random —Load-balance using packet random spray. |
| Required Privilege Level | routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Junos OS Routing Protocols and Policies Configuration Guide for Security Devices |

load-balancing-options (Aggregated Multiservices)

```
Syntax  load-balancing-options {
        high-availability-options {
            many-to-one {
                preferred-backup preferred-backup;
            }
        }
        member-failure-options {
            drop-member-traffic {
                rejoin-timeout rejoin-timeout;
            }
            redistribute-all-traffic {
                enable-rejoin;
            }
        }
        hash-keys {
            egress-key (destination-ip | source-ip);
            ingress-key (destination-ip | source-ip);
        }
        member-interface interface-name;
    }
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the high availability (HA) options for the aggregated Multiservices (AMS) interface.

Many-to-one (N:1) high availability mode for service applications like Network Address Translation (NAT) is supported. In this case, one Multiservices PIC is the backup (in hot standby mode) for one or more (N) active Multiservices PICs. If one of the active Multiservices PICs goes down, then the backup replaces it as the active Multiservices PIC. When the failed PIC comes back online, it becomes the new backup. This is called floating backup mode.

The remaining statements are explained separately.

Load-balancing might not be uniform among member interfaces in certain network deployments. The variance can be owing to a misconfiguration, which causes the traffic itself to be not sufficiently randomly distributed, causing the hash-keys to be ineffective. (for example, the hash key is destination IP but all sessions have only source IP address. The variation can be within the expected range and the load-balancing depends on the IP addresses chosen. The hash calculation performs a checksum on several bits of the IP address and not only on the last few lower significant bits of the IP address. In such a scenario, the load-balancing ratio can change, say, if the source IP address is changed from 20.0.0.0/24 to 20.0.1.0/24.

The distribution of traffic across member interfaces of an AMS interface is static load-balancing. Flows are load-balanced based on a packet-hash on parameters such as source IP or destination IP. Load-balancing effectiveness depends on the IP address

or protocol diversity. For example, if the hash-key is destination IP and all packets have the same destination, then all flows are directed to the same member. This is flow-level load-balancing and not per packet. As a result, traffic between a pair of addresses may be 10000 pps, whereas another pair of addresses may have 1 pps. The load of the former is not distributed among members. High availability is limited to stateless HA. When a backup interface takes over as an active interface, all flows are established afresh (for example, packets may undergo NAT processing differently after failover). However, such stateless failover does not impact other actives' running.

With a stateful firewall, static NAT as `basic-nat44` or `destination-nat44`, and dynamic NAT as `nat64`, `napt-44`, `dynamic-nat44`, and with application layer gateways (ALGs) configured, NAT hairpinning is not supported. Input direction for rule match to be applied is supported only for dynamic NAT types (NAT64, NAT44, and dynamic-NAT44). Service-set policies need to have "input" or "input-output" direction only. Flows on all active members are reset when the number of actives changes. The resetting of flows can be avoided at the cost of failed-member's traffic loss using certain options.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [interfaces on page 1386](#)
- [Understanding Aggregated Multiservices Interfaces on page 599](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 608](#)

local-certificate (Services IPsec VPN)

Syntax `local-certificate identifier;`

Hierarchy Level [edit [services](#) ipsec-vpn [ike policy](#) *policy-name*]

Release Information Statement introduced in Junos OS Release 7.5.

Description Name of the certificate that needs to be sent to the peer during the IKE authentication phase.

Options *identifier*—Name of certificate.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- [Configuring IKE Policies on page 409](#)

local-gateway (IPSec)

| | |
|---------------------------------|---|
| Syntax | <code>local-gateway <i>address</i>;</code> |
| Hierarchy Level | [edit services service-set service-set-name ipsec-vpn-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the local IPv4 or IPv6 address for the IPsec traffic. |
| Options | <i>address</i> —Local address. |
| Required Privilege Level | admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Service Rules on page 36 |

local-gateway (L2TP LNS)

| | |
|---------------------------------|---|
| Syntax | <pre>local-gateway { address <i>address</i>; gateway-name <i>gateway-name</i>; }</pre> |
| Hierarchy Level | [edit services l2tp tunnel-group name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the IP address or name for the local (LNS) gateway for L2TP tunnel.

The remaining statements are explained separately. |
| Options | <i>address</i> —Local IP address; corresponds to the IP address that is used by LACs to identify the LNS. When the LAC is an MX Series router, this address matches the remote gateway address configured in the LAC tunnel profile. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Local Gateway Address and PIC on page 642. • Configuring L2TP Tunnel Groups on page 641 • Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces |

local-id

| | |
|---------------------------------|---|
| Syntax | <code>local-id (ipv4_addr <i>ipv4-address</i> ipv6_addr <i>ipv6-address</i> key-id <i>identifier</i>);</code> |
| Hierarchy Level | [edit services ipsec-vpn ike policy <i>policy-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
<code>ipv6_addr</code> option added in Junos OS Release 7.6. |
| Description | Specify local identifiers for IKE Phase 1 negotiation. This statement is optional. |
| Options | <code>ipv4_addr <i>ipv4-address</i></code> —IPv4 address identification value.
<code>ipv6_addr <i>ipv6-address</i></code> —IPv6 address identification value.
<code>key_id <i>identifier</i></code> —Key identification value.
<code>fqdn <i>fqdn</i></code> —Fully-qualified domain name. |
| Required Privilege Level | <code>system</code> —To view this statement in the configuration.
<code>system-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Security Associations on page 385 |

log-prefix (L2TP)

| | |
|---------------------------------|---|
| Syntax | <code>log-prefix <i>prefix-value</i>;</code> |
| Hierarchy Level | [edit services l2tp tunnel-group <i>group-name</i> syslog host <i>hostname</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Set the system logging prefix value. |
| Options | <code><i>prefix-value</i></code> —System logging prefix value. |
| Required Privilege Level | <code>interface</code> —To view this statement in the configuration.
<code>interface-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring System Logging of L2TP Tunnel Activity on page 644 |

log-prefix (Services)

| | |
|---------------------------------|--|
| Syntax | <code>log-prefix <i>prefix-value</i>;</code> |
| Hierarchy Level | [edit services service-set service-set-name syslog host hostname] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Set the system logging prefix value. |
| Options | <i>prefix-value</i> —System logging prefix value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring System Logging for Service Sets on page 47 |

logging (Services)

| | |
|---------------------------------|--|
| Syntax | <pre>logging { traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; no-remote-trace; } }</pre> |
| Hierarchy Level | [edit services] |
| Release Information | Statement introduced in Junos OS Release 8.0. |
| Description | Define global services properties. |
| Options | The remaining statement is explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Tracing Services PIC Operations on page 48 |

logging (Services IDS)

| | |
|---------------------------------|---|
| Syntax | <pre>logging {
 syslog;
 threshold rate;
}</pre> |
| Hierarchy Level | [edit services ids rule rule-name term term-name then] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Set logging values for this IDS term. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Actions in IDS Rules on page 358 |

lsq-failure-options

| | |
|---------------------------------|--|
| Syntax | <pre>lsq-failure-options {
 no-termination-request;
 trigger-link-failure interface-name;
}</pre> |
| Hierarchy Level | [edit interfaces lsq-fpc/pic/port] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | For link services IQ (lsq) interfaces only, define the failure recovery option settings. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Association between LSQ and SONET Interfaces on page 546 |

manual

```
Syntax  manual {
        direction (inbound | outbound | bidirectional) {
            authentication {
                algorithm (hmac-md5-96 | hmac-sha1-96);
                key (ascii-text key | hexadecimal key);
            }
            auxiliary-spi spi-value;
            encryption {
                algorithm algorithm;
                key (ascii-text key | hexadecimal key);
            }
            spi spi-value;
            protocol (ah | esp | bundle);
        }
    }
```

Hierarchy Level [edit [services](#) ipsec-vpn [rule](#) *rule-name* [term](#) *term-name* [then](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define a manual IPsec SA.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring Security Associations on page 385](#)

many-to-one (Aggregated Multiservices)

| | |
|---------------------|--|
| Syntax | <pre>many-to-one {
 preferred-backup <i>preferred-backup</i>;
}</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> load-balancing-options high-availability-options] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Configure the initial preferred backup for the aggregated Multiservices (AMS) interface. |



NOTE: The preferred backup must be one of the member interfaces (*mams-*) that have already been configured at the [edit interfaces *interface-name* load-balancing-options] hierarchy level. Even in the case of mobile control plane redundancy, which is one-to-one (1:1), the initial preferred backup is configured at this hierarchy level.

The remaining statements are explained separately.

| | |
|--------------------------|--|
| Options | preferred-backup <i>preferred-backup</i> —Name of the preferred backup member interface.
The member interface format is mams-a/b/0 , where a is the Flexible PIC Concentrator (FPC) slot number and b is the PIC slot number. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• high-availability-options (aggregated Multiservices) on page 1376• Understanding Aggregated Multiservices Interfaces on page 599• Example: Configuring an Aggregated Multiservices Interface (AMS) on page 608 |

mapping-refresh

| | |
|---------------------------------|--|
| Syntax | mapping-refresh (inbound outbound inbound-outbound); |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated secure-nat-mapping] |
| Release Information | Statement introduced in Junos OS Release 12.3 |
| Description | Specify how the flow timer should be refreshed based on the mapping refresh configured for all types of fwnat flows. For TCP flows, if tcp-tickles is configured, then tickles are sent only on the flow matching the mapping-refresh direction. For inbound-outbound mapping, refresh tickles will be sent on both the flows (default behavior). |
| Options | <p>inbound—Refresh the flow timer for inbound flows only.</p> <p>inbound-outbound—Refresh the flow timer for all flows.</p> <p>outbound—Refresh the flow timer for outbound flows only.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Protecting CGN Devices Against Denial of Service (DOS) Attacks on page 241 |

mapping-timeout

| | |
|----------------------------|--|
| Syntax | mapping-timeout <i>seconds</i> ; |
| Hierarchy Level | [edit services nat pool <i>nat-pool-name</i>] |
| Release Information | mapping-timeout statement introduced in JUNOS Release 10.1. |



NOTE: This configuration option has been replaced by [app-mapping-timeout](#). This option is currently retained only for backward compatibility.

| | |
|---------------------------------|--|
| Description | Specify the duration for mappings that use the specified NAT pool. |
| Options | seconds —Lifetime of mappings in seconds.
Default: 300
Range: 120 through 864,000 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Source and Destination Addresses Network Address Translation Overview on page 66 |

match-direction (Services CoS)

| | |
|---------------------------------|--|
| Syntax | match-direction (input output input-output); |
| Hierarchy Level | [edit services cos rule <i>rule-name</i>] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | Specify the direction in which the rule match is applied. |
| Options | input —Apply the rule match on the input side of the interface.
output —Apply the rule match on the output side of the interface.
input-output —Apply the rule match bidirectionally. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring CoS Rules |

match-direction (Services IDS)

| | |
|---------------------------------|--|
| Syntax | match-direction (input output input-output); |
| Hierarchy Level | [edit services ids rule <i>rule-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the direction in which the rule match is applied. |
| Options | <p>input—Apply the rule match on input.</p> <p>output—Apply the rule match on output.</p> <p>input-output—Apply the rule match bidirectionally.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Match Conditions in IDS Rules on page 357 |

match-direction (Services IPsec VPN)

| | |
|---------------------------------|--|
| Syntax | match-direction (input output); |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the direction in which the rule match is applied. |
| Options | <p>input—Apply the rule match on input.</p> <p>output—Apply the rule match on output.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IPsec Rules on page 422 |

match-direction (Services NAT)

| | |
|---------------------------------|---|
| Syntax | match-direction (input output); |
| Hierarchy Level | [edit services nat rule <i>rule-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the direction in which the rule match is applied. |
| Options | input —Apply the rule match on input.
output —Apply the rule match on output. |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |


match-direction (Services Stateful Firewall)

| | |
|---------------------------------|--|
| Syntax | match-direction (input output input-output); |
| Hierarchy Level | [edit services stateful-firewall rule <i>rule-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the direction in which the rule match is applied. |
| Options | input —Apply the rule match on the input side of the interface.
output —Apply the rule match on the output side of the interface.
input-output —Apply the rule match bidirectionally. |
| Usage Guidelines | See “ Configuring Stateful Firewall Rules ” on page 331. |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |

max-drop-flows

| | |
|---------------------------------|---|
| Syntax | <pre>max-drop-flows { ingress <i>ingress-flows</i>; egress <i>egress-flows</i>; }</pre> |
| Hierarchy Level | [edit services service-set service-set-name] |
| Release Information | Statement introduced in Junos OS Release 12.3 |
| Description | <p>Configure the maximum drop flows allowed per ingress and egress direction. The configuration is per service set. The configured limits indicate the maximum number of drop flows that can be created at a given instance of time in both directions. If max drop flows ingress is 10 and egress is 5 then at a given instance of time maximum of 10 ingress drop flows and 5 egress drop flows can be present. Two counters, one for each direction ingress and egress, are to be added to service set stateful-firewall statistics to track the number of drop flows not created due to the drop flow limits exceeded. These limits applies to all types of drop flows i.e., TCP, UDP, ICMP etc. Ingress drop flows are forward flows for match-direction input rules and reverse flows for match-direction output rules. Similarly egress drop flows are reverse flows for match-direction input and forward flows for match-direction output rules. The limits are applied cumulatively on all the nat rules associated with the service-set.</p> |
| Options | <p><i>ingress-flows</i>—Maximum number of drop flows on the ingress interface.</p> <p><i>egress-flows</i>—Maximum number of drop flows on the egress interface.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Service Set Limitations on page 37 |


max-flows

| | |
|--|---|
| Syntax | <code>max-flows <i>number</i>;</code> |
| Hierarchy Level | [edit services service-set service-set-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Maximum number of flows allowed for the service set. |
| Options | <i>number</i> —Maximum number of flows. |
| <hr/> | |
| <div> NOTE: When an aggregated multiservices (AMS) interface is configured as the service interface for a service set, the max-flow value configured for the service set is applied to each of the member interfaces in the AMS interface. That is, if you have configured 1000 as the max-flow value for a service set that uses an AMS interface with four active member interfaces, each of the member interfaces can handle 1000 flows each, resulting in an effective max-flow value of 4000.</div> <hr/> | |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Set Limitations on page 37 |

maximum-contexts

| | |
|---------------------------------|---|
| Syntax | <code>maximum-contexts <i>number</i> <force>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression <i>rtp</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression <i>rtp</i>] |
| Release Information | Statement introduced in Junos OS Release 7.5. |
| Description | Specify the maximum number of RTP contexts to accept during negotiation. |
| Options | <p><i>number</i>—Maximum number of contexts.</p> <p><i>force</i>—(Optional) Requires the PIC to use the value specified for maximum RTP contexts, regardless of the negotiated value. This option allows the software to interoperate with Junos OS Releases that base the RTP context value on link speed.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Compression of Voice Traffic on page 623 |

maximum-send-window

| | |
|--|--|
| Syntax | <code>maximum-send-window <i>packets</i>;</code> |
| Hierarchy Level | [edit services l2tp tunnel-group <i>name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the size of the send window for L2TP tunnels, which limits the remote end's receive window size. |
| <div style="display: flex; align-items: center;">  <div> <p>NOTE: This statement is not supported for L2TP LNS on MX Series routers.</p> </div> </div> | |
| Options | <p><i>packets</i>—Maximum number of packets the send window can hold at one time.</p> <p>Default: 32</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Window Size for L2TP Tunnels on page 643 |

member-failure-options (Aggregated Multiservices)

Syntax

```
member-failure-options {
  drop-member-traffic {
    rejoin-timeout rejoin-timeout;
  }
  redistribute-all-traffic {
    enable-rejoin;
  }
}
```

Hierarchy Level [edit interfaces *interface-name* load-balancing-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the possible behavior for the aggregated Multiservices (AMS) interface in case of failure of more than one active member.



NOTE: The `drop-member-traffic` configuration and the `redistribute-all-traffic` configuration are mutually exclusive.

Table 52 on page 1408 displays the behavior of the member interface after the failure of the first Multiservices PIC. Table 53 on page 1409 displays the behavior of the member interface after the failure of two Multiservices PICs.



NOTE: The AMS infrastructure has been designed to handle one failure automatically. However, in the unlikely event that more than one Multiservices PIC fails, the AMS infrastructure provides configuration options to minimize the impact on existing traffic flows.

Table 52: Behavior of Member Interface After One Multiservices PIC Fails

| High Availability Mode | Member Interface Behavior |
|--|---|
| Many-to-one (N:1) high availability support for service applications | Automatically handled by the AMS infrastructure |

Table 53: Behavior of Member Interface After Two Multiservices PICs Fail

| High Availability Mode | Configuration | rejoin-timeout | Behavior when member rejoins before rejoin-timeout expires | Behavior when member rejoins after rejoin-timeout expires |
|--|---------------------------------|----------------|---|---|
| Many-to-one (N:1) high availability support for service applications | drop-member-traffic | Configured | <p>The existing traffic for the second failed member will <i>not</i> be redistributed to the other members.</p> <p>The first member to rejoin becomes an active member. The second member to rejoin becomes the backup. This behavior is handled automatically by the AMS infrastructure.</p> | <p>The existing traffic for the second failed member will <i>not</i> be redistributed to the other members.</p> <p>The first member will rejoin the AMS automatically. However, the other members who are rejoining will be moved to the discard state.</p> |
| Many-to-one (N:1) high availability support for service applications | redistribute-all-traffic | Not applicable | <p>Before rejoin, the traffic is redistributed to existing active members.</p> <p>After a failed member rejoins, the traffic is load-balanced afresh. This may impact existing traffic flows.</p> | |

The remaining statements are explained separately.


Default If **member-failure-options** are not configured, then the default behavior is to drop member traffic with a rejoin timeout of 120 seconds.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.


Related Documentation

- [load-balancing-options \(Aggregated Multiservices\) on page 1393](#)
- [Understanding Aggregated Multiservices Interfaces on page 599](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 608](#)

member-interface (Aggregated Multiservices)

| | |
|---------------------------------|--|
| Syntax | <code>member-interface <i>interface-name</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> load-balancing-options] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | <p>Specify the member interfaces for the aggregated Multiservices (AMS) interface. You can configure multiple interfaces by specifying each interface in a separate statement.</p> <p>For high availability service applications like Network Address Translation (NAT) that support many-to-one (N:1) redundancy, you can specify two or more interfaces.</p> |
| | <div> NOTE: The member interfaces that you specify must be members of aggregated Multiservices interfaces (mams-).</div> |
| | <p>The remaining statements are explained separately.</p> |
| Options | <p><i>interface-name</i>—Name of the member interface. The member interface format is mams-a/b/0, where a is the Flexible PIC Concentrator (FPC) slot number and b is the PIC slot number.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Understanding Aggregated Multiservices Interfaces on page 599• Example: Configuring an Aggregated Multiservices Interface (AMS) on page 608• load-balancing-options (Aggregated Multiservices) on page 1393 |

message-rate-limit

| | |
|--|---|
| Syntax | <code>message-rate-limit <i>messages-per-second</i></code> |
| Hierarchy Level | <pre> interfaces <i>interface-name</i> { services-options { <i>cg</i>n-pic; disable-global-timeout-override; ignore-errors <<i>alg</i>> <tcp>; inactivity-non-tcp-timeout <i>seconds</i>; inactivity-tcp-timeout <i>seconds</i>; inactivity-timeout <i>seconds</i>; open-timeout <i>seconds</i>; session-limit { maximum <i>number</i>; rate <i>new-sessions-per-second</i>; } session-timeout <i>seconds</i>; syslog { } } }</pre> |
| Release Information | Statement introduced Junos OS Release 11.1. |
| Description | Maximum system log messages per second allowed from this interface. |
| <div>  <p>NOTE: The message-rate-limit command can be configured only for physical service interfaces (sp-x/x/x) and not for redundancy services PIC interfaces (rspx).</p> </div> | |
| Options | <p><i>messages-per-second</i>—This option configures the maximum number of system log messages per second that can be formatted and sent from the PIC to either the Routing Engine (local) or to an external server (remote). The default rates are 10,000 for the Routing Engine and 800,000 for an external server.</p> <p>Range: 0 through 2147483647</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> Configuring System Logging for Service Sets on page 47 |

mlfr-uni-nni-bundles-inline

| | |
|---------------------------------|---|
| Syntax | mlfr-uni-nni-bundles-inline <i>number</i> ; |
| Hierarchy Level | [edit chassis fpc <i>number</i> pic <i>number</i>] |
| Release Information | Statement introduced in Junos OS Release 14.1. |
| Description | Specify the number of inline multilink frame relay user-to-network interface (UNI) and network-to-network interface (NNI) bundles. |
| Options | <i>number</i> —Specify the number of inline multilink frame relay UNI NNI bundles.
Range: 1 through 255 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Junos OS to Support the Link Services PIC• Inline MLPPP for WAN Interfaces Overview on page 559• Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces on page 754• Example: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces on page 788 |

mode (Services IPsec VPN)

| | |
|---------------------------------|--|
| Syntax | <code>mode (aggressive main);</code> |
| Hierarchy Level | [edit services ipsec-vpn ike policy <i>policy-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define an IKE policy mode. |
| Default | <code>main</code> |
| Options | <p>aggressive—Takes half the number of messages of main mode, has less negotiation power, and does not provide identity protection.</p> <p>main—Uses six messages, in three peer-to-peer exchanges, to establish the IKE SA. These three steps include the IKE SA negotiation, a Diffie-Hellman exchange, and authentication of the peer. Also provides identity protection.</p> |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IKE Policies on page 409 |

mss

| | |
|---------------------------------|---|
| Syntax | <code>mss <i>value</i>;</code> |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> then <i>syn-cookie</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the maximum segment size (MSS) value used in Transmission Control Protocol (TCP) delayed binding. |
| Options | <p>value—MSS value.</p> <p>Default: 1500</p> <p>Range: 128 through 8192</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Actions in IDS Rules on page 358 |

multi-link-layer-2-inline

| | |
|---------------------------------|---|
| Syntax | multi-link-layer-2-inline; |
| Hierarchy Level | [edit chassis fpc <i>number</i> pic <i>number</i>] |
| Release Information | Statement introduced in Junos OS Release 14.1. |
| Description | Enable inline Layer 2 bundling services. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Junos OS to Support the Link Services PIC• Inline MLPPP for WAN Interfaces Overview on page 559• Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces on page 754• Example: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces on page 788 |

multilink-class

| | |
|---------------------------------|--|
| Syntax | multilink-class <i>number</i> ; |
| Hierarchy Level | [edit class-of-service fragmentation-maps <i>map-name</i> forwarding-class <i>class-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services IQ (lsq) interfaces only, map a forwarding class into a multiclass MLPPP (MCML).

The multilink-class statement and no-fragmentation statements are mutually exclusive. |
| Options | number —The multilink class assigned to this forwarding class.
Range: 0 through 7
Default: None |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 527• Configuring Multiclass MLPPP on LSQ Interfaces on page 562• Configuring Fragmentation by Forwarding Class• Junos OS Services Interfaces Library for Routing Devices• multilink-max-classes on page 1415 |

multilink-max-classes

| | |
|---------------------------------|--|
| Syntax | <code>multilink-max-classes <i>number</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services IQ (lsq) interfaces only, configure the number of multilink classes to be negotiated when a link joins the bundle. |
| Options | <i>number</i> —The number of multilink classes to be negotiated when a link joins the bundle.
Range: 1 through 8
Default: None |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Multiclass MLPPP on LSQ Interfaces on page 562 |


nat-options

| | |
|---------------------------------|---|
| Syntax | <pre> nat-options { land-attack-check (ip-only ip-port); max-sessions-per-subscriber <i>session-number</i>; stateful-nat64 { clear-dont-fragment-bit; } } </pre> |
| Hierarchy Level | [edit services service-set <i>service-set-name</i>] |
| Release Information | Statement introduced with Junos OS Release 12.1.
land-attack-check and max-sessions-per-subscriber statements added in 13.3. |
| Description | Specify parameters for NAT operation. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Service Rules on page 36 • <i>clear-dont-fragment-bit</i> • <i>land-attack-check</i> • <i>max-sessions-per-subscriber</i> • <i>stateful-nat64</i> |

nat-rules

| | |
|---------------------------------|---|
| Syntax | (nat-rules <i>rule-name</i> nat-rule-sets <i>rule-set-name</i>); |
| Hierarchy Level | [edit services service-set service-set-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the Network Address Translation (NAT) rules or rule set included in this service set. You can configure multiple rules, but only one rule set for each service. |
| Options | <i>rule-name</i> —Identifier for the collection of terms that constitute this rule.
<i>rule-set-name</i> —Identifier for the set of rules to be included. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Rules on page 36 |

next-hop-service

| | |
|---|--|
| Syntax | <pre> next-hop-service { inside-service-interface <i>interface-name.unit-number</i>; outside-service-interface <i>interface-name.unit-number</i>; outside-service-interface-type local; service-interface-pool <i>name</i>; } </pre> |
| Hierarchy Level | [edit services service-set service-set-name] |
| Release Information | Statement introduced before Junos OS Release 7.4.
service-interface-pool option added in Junos OS Release 9.3. |
| Description | Specify interface names or a service interface pool for the forwarding next-hop service set. You cannot specify both a service interface pool and an inside or outside interface. |
| Options | <p>inside-service-interface <i>interface-name.unit-number</i>—Name and logical unit number of the service interface associated with the service set applied inside the network.</p> <p>outside-service-interface <i>interface-name.unit-number</i>—Name and logical unit number of the service interface associated with the service set applied outside the network.</p> <p>outside-service-interface-type <i>interface-type</i>—Identifies the interface type of the service interface associated with the service set applied outside the network. For inline IP reassembly, set the interface type to local.</p> <p>service-interface-pool <i>name</i>—Name of the pool of logical interfaces configured at the [edit services service-interface-pools pool <i>pool-name</i>] hierarchy level. You can configure a service interface pool only if the service set has a PGCP rule configured. The service set cannot contain any other type of rule.</p> |
| <div>  <p>NOTE: service-interface-pool is not applicable for IP reassembly configuration on L2TP.</p> </div> | |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> Configuring Service Sets to be Applied to Services Interfaces on page 31 |

no-anti-replay (Services IPsec VPN)

| | |
|--------------------------|---|
| Syntax | no-anti-replay; |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Disable IPsec antireplay service, which occasionally causes interoperability issues for security associations. |
| Required Privilege Level | admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Security Associations on page 385 |

no-anti-replay (Services Service Set)

| | |
|---------------------|--|
| Syntax | no-anti-replay; |
| Hierarchy Level | [edit services service-set <i>service-set-name</i> ipsec-vpn-options] |
| Release Information | Statement introduced in Junos OS Release 10.0. |
| Description | <p>Disable IPsec antireplay service for this service set, which occasionally causes interoperability issues for security associations. This statement is useful for dynamic endpoint tunnels for which you cannot configure the no-anti-reply statement at the [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] hierarchy level.</p> <p>For static IPsec tunnels, this statement disables the antireplay check for all the tunnels within this service set. If antireplay check has to be enabled for a particular tunnel, then set the anti-replay-window-size statement at the [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] hierarchy level.</p> |



NOTE: Setting the anti-replay-window-size and no-anti-replay statements at the [edit [services](#) ipsec-vpn rule *rule-name* term *term-name* then] hierarchy level overrides the settings specified at the [edit [services](#) service-set *service-set-name* ipsec-vpn-options] hierarchy level.

| | |
|--------------------------|---|
| Usage Guidelines | See or . |
| Required Privilege Level | admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Service Sets on page 430• Configuring or Disabling IPsec Anti-Replay |

no-fragmentation

| | |
|---------------------------------|--|
| Syntax | no-fragmentation; |
| Hierarchy Level | [edit class-of-service fragmentation-maps forwarding-class <i>class-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>For link services IQ (lsq) interfaces only, set traffic on a particular forwarding class to be interleaved, rather than fragmented. This statement specifies that no extra fragmentation header is prepended to the packets received on this queue and that static-link load balancing is used to ensure in-order packet delivery.</p> <p>Static-link load balancing is done based on packet payload. For IP version 4 (IPv4) and IP version 6 (IPv6) traffic, the link is chosen based on a hash computed from the source address, destination address, and protocol. If the IP payload is Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) traffic, the hash also includes source port and destination port. For MPLS traffic, the hash includes all MPLS labels and fields in the payload, whether the MPLS payload is IPv4 or IPv6.</p> |
| Default | If you do not include this statement, the traffic in forwarding class <i>class-name</i> is fragmented. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 527 |

no-ipsec-tunnel-in-traceroute

| | |
|---------------------------------|--|
| Syntax | no-ipsec-tunnel-in-traceroute; |
| Hierarchy Level | [edit services ipsec-vpn] |
| Release Information | Statement introduced in Junos OS Release 10.0. |
| Description | <p>Disables displaying the IPsec tunnel endpoint in the trace route output. The IPsec tunnel is not treated as a next hop and TTL is not decremented. If the TTL becomes zero, the ICMP time exceeded message will not be generated.</p> |
| Required Privilege Level | admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Security Associations on page 385 |

no-per-unit-scheduler

| | |
|---------------------------------|---|
| Syntax | no-per-unit-scheduler; |
| Hierarchy Level | [edit interfaces <i>interface-name</i>] |
| Release Information | Statement introduced before Junos OS Release 11.4. |
| Description | To enable traffic control profiles to be applied at FRF.16 bundle (physical) interface level, disable the per-unit scheduler, which is enabled by default. This statement and the shared-scheduler statement are mutually exclusive. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Oversubscribing Interface Bandwidth |

no-termination-request

| | |
|---------------------------------|--|
| Syntax | no-termination-request; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> ppp-options],
[edit interfaces <i>lsq-fpc/pic/port</i> <i>lsq-failure-options</i>] |
| Release Information | Statement introduced in Junos OS Release 7.4.
Support at the [edit interfaces <i>interface-name</i> ppp-options] hierarchy level added in Junos OS Release 8.3. |
| Description | Inhibit PPP termination-request messages to the remote host if the primary circuit fails. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Association between LSQ and SONET Interfaces on page 546 |

no-translation

| | |
|---------------------------------|---|
| Syntax | no-translation; |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced in Junos OS Release 7.6. |
| Description | Specify that traffic is not to be translated. |
| Options | none |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Network Address Translation Rules Overview on page 69 |

output

| | |
|---------------------------------|---|
| Syntax | output {
[service-set <i>service-set-name</i> < service-filter <i>filter-name</i> >];
} |
| Hierarchy Level | [edit interface <i>interface-name</i> unit <i>logical-unit-number</i> family inet service],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the output service sets and filters to be applied to traffic. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Applying Filters and Services to Interfaces on page 38 |


overload-pool

| | |
|---------------------------------|--|
| Syntax | <code>overload-pool <i>overload-pool-name</i>;</code> |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated] |
| Release Information | Statement introduced in Junos OS Release 7.6. |
| Description | Specify an address pool that can be used if the source pool becomes exhausted. |
| Options | <i>overload-pool-name</i> —Name of the overload pool. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |

overload-prefix

| | |
|---------------------------------|--|
| Syntax | <code>overload-prefix <i>overload-prefix</i>;</code> |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated] |
| Release Information | Statement introduced in Junos OS Release 7.6. |
| Description | Specify the prefix that can be used if the source pool becomes exhausted. |
| Options | <i>overload-prefix</i> —Prefix value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |

passive-mode-tunneling

| | |
|---------------------------------|--|
| Syntax | <code>passive-mode-tunneling;</code> |
| Hierarchy Level | [edit services service-set service-set-name ipsec-vpn-options] |
| Release Information | Statement introduced in Junos OS Release 10.0. |
| Description | <p>Allows tunneling of malformed packets. When this statement is enabled, traffic bypasses the usual active IP checks. The IPsec tunnel is not treated as a next hop and TTL is not decremented. If the packet size exceeds the tunnel MTU value, an ICMP error is not generated. Starting with Junos OS Release 13.3R4 and 14.2R1, passive mode tunneling is supported on MS-MICs and MS-MPCs.</p> <hr/> <div>  <p>NOTE: The header-integrity-check option that is supported on MS-MICs and MS-MPCs to verify the packet header for anomalies in IP, TCP, UDP, and ICMP information and flag such anomalies and errors has a functionality that is opposite to the functionality caused by passive mode tunneling. If you configure both the header-integrity-check statement and the passive-mode tunneling statement on MS-MICs and MS-MPCs, and attempt to commit such a configuration, an error is displayed during commit.</p> <p>The passive mode tunneling functionality (by including the <code>passive-mode-tunneling</code> statement at the [edit <code>services service-set service-set-name ipsec-vpn-options</code>] hierarchy level) is a superset of the capability to disable IPsec tunnel endpoint in the traceroute output (by including <code>no-ipsec-tunnel-in-traceroute</code> statement at the [edit <code>services ipsec-vpn</code>] hierarchy level). Passive mode tunneling also bypasses the active IP checks and tunnel MTU check in addition to not treating an IPsec tunnel as a next-hop as configured by the <code>no-ipsec-tunnel-in-traceroute</code> statement.</p> <hr/> </div> |
| Required Privilege Level | admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IPsec Service Sets on page 430 |

per-unit-scheduler

| | |
|----------------------------|---|
| Syntax | <code>per-unit-scheduler;</code> |
| Hierarchy Level | <code>[edit interfaces <i>interface-name</i>]</code> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 13.2 on 16x10GE MPC and MPC3E line cards.</p> <p>Statement introduced in Junos OS Release 13.2 on PTX Series Packet Transport Routers.</p> <p>Statement introduced in Junos OS Release 13.3 on MPC4E line cards.</p> <p>Statement introduced in Junos OS Release 15.1 on MPC6E line cards.</p> |
| Description | For Channelized OC3 IQ, Channelized OC12 IQ, Channelized STM1 IQ, Channelized T3 IQ, Channelized E1 IQ, E3 IQ, link services IQ interfaces (lsq-), Gigabit Ethernet IQ, Gigabit Ethernet IQ2 and IQ2-E, and 10-, 40-, and 100-Gigabit Ethernet interfaces (including the 16x10GE MPC), enable the association of scheduler map names with logical interfaces. |



CAUTION: Turning on per-unit scheduling causes the interface to reinitialize, which means all logical interfaces (units) on the interface are deleted and recreated.



NOTE: Per-unit scheduling is not supported on T1 interfaces configured on the Channelized OC12 IQ PIC.



NOTE: On Gigabit Ethernet IQ2 and IQ2-E PICs without the `per-unit-scheduler` statement, the entire PIC supports 4071 VLANs and the user can configure all the VLANs on the same port.

On Gigabit Ethernet IQ2 and IQ2-E PICs with the `per-unit-scheduler` statement, the entire PIC supports $1024 - 2 * \text{number of ports}$ (1024 minus two times the number of ports), because each port is allocated two default schedulers.

When including the `per-unit-scheduler` statement, you must also include the `vlan-tagging` statement or the `flexible-vlan-tagging` statement (to apply scheduling to VLANs) or the `encapsulation frame-relay` statement (to apply scheduling to DLCIs) at the `[edit interfaces interface-name]` hierarchy level.

| | |
|---------------------------------|--|
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> <i>Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs</i> |

- *vlan-tagging*
- *flexible-vlan-tagging*
- *Example: Applying Scheduling and Shaping to VLANs*
- *Configuring Virtual LAN Queuing and Shaping on PTX Series Routers*

perfect-forward-secrecy (Services IPsec VPN)

| | |
|---------------------------------|--|
| Syntax | perfect-forward-secrecy {
keys (group1 group2 group5 group14 group19 group20);
} |
| Hierarchy Level | [edit services ipsec-vpn ipsec policy <i>policy-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define Perfect Forward Secrecy (PFS). Creates single-use keys. This statement is optional. |
| Options | <p>keys—Type of Diffie-Hellman prime modulus group that IKE uses when performing the new Diffie-Hellman exchange. The key can be one of the following:</p> <ul style="list-style-type: none"> • group1—768-bit. • group2—1024-bit. • group5—1536-bit. • group14—2048-bit. • group19—256-bit random Elliptic Curve Group. • group20—384-bit random Elliptic Curve Group. |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Security Associations on page 385 |

pgcp-rules

| | |
|---------------------------------|---|
| Syntax | (pgcp-rules <i>rule-name</i> pgcp-rules-sets <i>rule-set-name</i>); |
| Hierarchy Level | [edit services service-set service-set-name] |
| Release Information | Statement introduced in Junos OS Release 8.4. |
| Description | Specify the Packet Gateway Control Protocol (PGCP) rules or rule set included in this service set. You can configure multiple rules but only one rule set for each service. |
| Options | <i>rule-name</i> —Identifier for the collection of terms that constitute this rule.
<i>rule-set-name</i> —Identifier for the set of rules to be included. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface—control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Sets to be Applied to Services Interfaces on page 31 |

policy (Services IKE)

Syntax `policy policy-name {
 description description;
 local-certificate identifier;
 local-id (ipv4_addr ipv4-address | ipv6_addr ipv6-address | key-id identifier);
 version (1 | 2);
 mode (aggressive | main);
 pre-shared-key (ascii-text key | hexadecimal key);
 proposals [proposal-names];
 remote-id {
 any-remote-id;
 ipv4_addr [values];
 ipv6_addr [values];
 key_id [values];
 }
 respond-bad-spi max-responses
 }`

Hierarchy Level [edit [services](#) ipsec-vpn [ike](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define an IKE policy.

Options *policy-name*—IKE policy name.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
 admin-control—To add this statement to the configuration.

Related Documentation • [Configuring IKE Policies on page 409](#)

policy (Services IPsec VPN)

| | |
|---------------------------------|---|
| Syntax | <pre>policy <i>policy-name</i> {
 <i>description</i> <i>description</i>;
 perfect-forward-secrecy {
 keys (group1 group 14 group2 group 5);
 }
 proposals [<i>proposal-names</i>];
}</pre> |
| Hierarchy Level | [edit services ipsec-vpn ipsec] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define an IPsec policy. |
| Options | <p><i>policy-name</i>—IPsec policy name.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Policies on page 420 |

pool

Syntax `pool nat-pool-name {`
 `address ip-prefix </prefix-length>;`
 `address-allocation round-robin;`
 `address-range low minimum-value high maximum-value;`
 `app-mapping-timeout app-mapping-timeout;`
 `ei-mapping-timeout ei-mapping-timeout;`
 `mapping-timeout mapping-timeout;`
 `pgcp {`
 `hint [hint-strings];`
 `ports-per-session ports;`
 `remotely-controlled;`
 `}`
 `port {`
 `automatic (sequential | random-allocation);`
 `range low minimum-value high maximum-value random-allocation;`
 `preserve-parity;`
 `preserve-range;`
 `secured-port-block-allocation {`
 `active-block-timeout timeout-seconds;`
 `block-size block-size;`
 `max-blocks-per-user max-blocks;`
 `}`
 `}`
`}`

Hierarchy Level [edit [services](#) nat]

Release Information Statement introduced before Junos OS Release 7.4.
pgcp statement added in Junos OS Release 8.4.
remotely-controlled and **ports-per-session** statements added in Junos OS Release 8.5.
hint statement added in Junos OS Release 9.0.
address-allocation statement added in Junos OS Release 11.2.
sequential statement introduced in Junos OS Release 14.2.

Description Specify the NAT name and properties.

Options *nat-pool-name*—Identifier for the NAT address pool.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Pools of Addresses and Ports for Network Address Translation Overview on page 67](#)

port (Services NAT)

Syntax port {
 automatic (sequential | random-allocation);
 range low *minimum-value* high *maximum-value* random-allocation;
 preserve-parity;
 preserve-range;
 deterministic-port-block-allocation <block-size *block-size*> <include-boundary-addresses>;
 secured-port-block-allocation {
 active-block-timeout *timeout-seconds*;
 block-size *block-size*;
 max-blocks-per-user *max-blocks*;
 }
 }
}

Hierarchy Level [edit [services](#) nat [pool](#) *nat-pool-name*]

Release Information port statement introduced before Junos OS Release 7.4.
 random-allocation statement introduced in Junos OS Release 9.3.
 secured-port-block-allocation statement introduced in Junos OS Release 11.2.
 deterministic-port-block-allocation statement introduced in Junos OS Release 12.1.
 sequential statement introduced in Junos OS Release 14.2.

Description Specify the NAT pool port or range. You can configure an automatically assigned port or specify a range with minimum and maximum values.



NOTE: Until Junos OS release 14.1, you could include the port automatic statement at the [edit services nat pool *nat-pool-name*] hierarchy level without having to use the auto option with the port automatic statement. Although the default method of assignment of ports was sequential (indicated by the auto option), the auto option was not required to be specified. Starting with Junos OS release 14.2, the sequential option is introduced to enable you to configure sequential allocation of ports. The sequential and random-allocation options available with the port automatic statement at the [edit services nat pool *nat-pool-name*] hierarchy level are mutually exclusive. You can include the sequential option for sequential allocation and the random-allocation option for random delegation of ports. By default, sequential allocation of ports takes place if you include only the port automatic statement at the [edit services nat pool *nat-pool-name*] hierarchy level. The auto option is hidden and is deprecated in Junos OS Release 14.2 and later, and is only maintained for backward compatibility. It might be removed completely in a future software release.

If you upgrade a router running a Junos OS release earlier than Release 14.2 to Release 14.2 and if the router contains the port automatic statement defined without the auto option included with the configuration, the router validates the auto option present in the configuration for sequential allocation of ports.

Options **automatic**—Cause the port assignment type to be automatically performed by the router.

sequential—Allocate ports in a sequential manner. With sequential allocation, the next available address in the NAT pool is selected only when all the ports available from an address are exhausted.



NOTE: The **auto** option is hidden and is deprecated in Junos OS Release 14.2 and later, and is only maintained for backward compatibility. It might be removed completely in a future software release. Starting with Junos OS Release 14.2, you must use the **sequential** option to allocate ports in a sequential manner, which is the default mode of allocation of ports.

minimum-value—Lower boundary for the port range.

maximum-value—Upper boundary for the port range.

preserve-parity—Allocate ports with same parity as the original port.

preserve-range—Preserve privileged port range after translation.

random-allocation—Allocate ports within a specified range randomly.

Other options are described separately.

Required Privilege **interface**—To view this statement in the configuration.

Level **interface-control**—To add this statement to the configuration.

- Related Documentation**
- [Configuring Source and Destination Addresses Network Address Translation Overview on page 66](#)
 - [Configuring Address Pools for Network Address Port Translation \(NAPT\) Overview on page 117](#)

port (Services Voice)

| | |
|---------------------------------|--|
| Syntax | <pre>port {
 minimum <i>port-number</i>;
 maximum <i>port-number</i>;
}</pre> |
| Hierarchy Level | [edit interfaces <i>lsq-fpc/pic/port</i> unit <i>logical-unit-number</i> compression <i>rtp</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>lsq-fpc/pic/port</i> unit <i>logical-unit-number</i> compression <i>rtp</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For voice services interfaces only, specify a range of User Datagram Protocol (UDP) destination port numbers in which RTP compression takes place. |
| Options | minimum <i>port-number</i> —Specify the minimum port number.
Range: 0 through 65,535

maximum <i>port-number</i> —Specify the maximum port number.
Range: 0 through 65,535 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Compression of Voice Traffic on page 623 |

port (System Log Messages)

| | |
|---------------------------------|---|
| Syntax | <pre>port <i>port-number</i>;</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options syslog host <i>hostname</i>] |
| Release Information | Statement introduced in Junos OS Release 11.1. |
| Description | UDP port for system log messages on the host. The default port is 514. |
| Options | <i>port-number</i> —Port number for system log messages. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring System Logging for Services Interfaces on page 20 |

port-forwarding

| | |
|---------------------------------|---|
| Syntax | <code>port-forwarding <i>map-name</i> {
 <i>destined-port</i>;
 <i>translated-port</i>;
}</code> |
| Hierarchy Level | [edit services nat] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Specify the mapping for port forwarding. |
| Options | <i>map-name</i> —Identifier for the port forwarding map. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Port Forwarding for Static Destination Address Translation on page 183 • Configuring Port Forwarding Without Destination Address Translation on page 186 |

port-forwarding-mappings

| | |
|---------------------------------|---|
| Syntax | <code>port-forwarding-mappings <i>map-name</i>;</code> |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Specify the name for mapping port forwarding in a Network Address Translation configuration. |
| Options | <i>map-name</i> —Identifier for the port forwarding mapping. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Port Forwarding for Static Destination Address Translation on page 183 • Configuring Port Forwarding Without Destination Address Translation on page 186 |

ports-per-session

| | |
|---------------------------------|--|
| Syntax | <code>ports-per-session <i>ports</i>;</code> |
| Hierarchy Level | [edit services nat pool <i>nat-pool-name</i> pgcp] |
| Release Information | Statement introduced in Junos OS Release 8.4. |
| Description | Configure the number of ports required to support Real-Time Transport Protocol (RTP), Real-Time Control Protocol (RTCP), Real-Time Streaming Protocol (RTSP), and forward error correction (FEC) for voice and video flows on the Multiservices PIC. |
| Options | <i>number-of-ports</i> —Number of ports to enable: 2 or 4 for combined voice and video services.
Default: 2 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

post-service-filter

| | |
|---------------------------------|--|
| Syntax | <code>post-service-filter <i>filter-name</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service input],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service input] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the filter to be applied to traffic after service processing. The filter is applied only if a service set is configured and selected. You can configure a postservice filter on the input side of the interface only.

The post-service-filter statement is not supported when the service interface is on an MS-MIC or MS-MPC. |
| Options | <i>filter-name</i> —Identifier for the post-service filter. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Applying Filters and Services to Interfaces on page 38 |

ppp-access-profile

| | |
|----------------------------|--|
| Syntax | <code>ppp-access-profile <i>profile-name</i>;</code> |
| Hierarchy Level | [edit services l2tp tunnel-group <i>name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the profile used to validate all Point-to-Point Protocol (PPP) session requests through L2TP tunnels established to the local gateway address. |



NOTE: This statement is not supported for L2TP LNS on MX Series routers.

| | |
|---------------------------------|--|
| Options | <i>profile-name</i> —Identifier for the PPP profile. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Access Profiles for L2TP Tunnel Groups on page 642 |

pre-shared-key (Services IKE)

| | |
|---------------------------------|---|
| Syntax | <code>pre-shared-key (ascii-text <i>key</i> hexadecimal <i>key</i>);</code> |
| Hierarchy Level | [edit services ike policy <i>policy-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define a preshared key for an IKE policy. |
| Options | <i>key</i> —Value of preshared key. The key can be one of the following: <ul style="list-style-type: none"> • ascii-text—ASCII text key. • hexadecimal—Hexadecimal key. |
| Required Privilege Level | admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IKE Policies on page 409 |

preserve-interface

| | |
|---------------------------------|--|
| Syntax | <code>preserve-interface;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> sonet-options aps] |
| Release Information | Statement introduced in Junos OS Release 7.6. |
| Description | <p>Provide link PIC replication, providing MLPPP link redundancy at the port level. This feature is supported with SONET APS and the following link PICs:</p> <ul style="list-style-type: none">• Channelized OC3 IQ PIC• Channelized OC12 IQ PIC• Channelized STM1 IQ PIC <p>Link PIC replication provides the ability to add two sets of links, one from the active SONET PIC and the other from the standby SONET PIC, to the same bundle. If the active SONET PIC fails, links from the standby PIC are used without triggering link renegotiation. All the negotiated state is replicated from the active links to the standby links to prevent link renegotiation.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Link State Replication for Redundant Link PICs on page 551 |

primary (Adaptive Services Interfaces)

| | |
|---------------------------------|---|
| Syntax | <code>primary interface-name;</code> |
| Hierarchy Level | [edit interfaces (rsp0 rsp1) redundancy-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the primary adaptive services interface. |
| Options | <i>interface-name</i> —The identifier for the AS or Multiservices PIC interface, which must be of the form <i>sp-fpc/pic/port</i> . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring AS or Multiservices PIC Redundancy on page 41 |

primary (Link Services IQ PIC Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>primary interface-name;</code> |
| Hierarchy Level | [edit interfaces <code>rlsnumber</code> <code>redundancy-options</code>] |
| Release Information | Statement introduced in Junos OS Release 7.6. |
| Description | Specify the primary Link Services IQ PIC interface. |
| Options | <i>interface-name</i> —The identifier for the Link Services IQ PIC interface, which must be of the form <code>lsq-fpc/pic/port</code> . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 548 |

proposal (Services IKE)

| | |
|---------------------------------|---|
| Syntax | <pre>proposal proposal-name { authentication-algorithm (md5 sha1 sha-256); authentication-method (dsa-signatures pre-shared-keys rsa-signatures); description description; dh-group (group1 group2 group5 group14); encryption-algorithm algorithm; lifetime-seconds seconds; }</pre> |
| Hierarchy Level | [edit <code>services ipsec-vpn ike</code>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define an IKE proposal for a dynamic SA. |
| Options | <i>proposal-name</i> —IKE proposal name.

The remaining statements are explained separately. |
| Required Privilege Level | system—To view this statement in the configuration.
system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IKE Proposals on page 405 |

proposal (Services IPsec VPN)

| | |
|---------------------------------|--|
| Syntax | <pre>proposal <i>proposal-name</i> {
 authentication-algorithm (hmac-md5-96 hmac-sha1-96);
 description <i>description</i>;
 encryption-algorithm <i>algorithm</i>;
 lifetime-seconds <i>seconds</i>;
 protocol (ah esp bundle);
}</pre> |
| Hierarchy Level | [edit services ipsec-vpn ipsec] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define an IPsec proposal for a dynamic SA. |
| Options | <p><i>proposal-name</i>—IPsec proposal name.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Proposals on page 415 |

proposals

| | |
|---------------------------------|---|
| Syntax | <pre>proposals [<i>proposal-names</i>];</pre> |
| Hierarchy Level | [edit services ipsec-vpn ike policy <i>policy-name</i>],
[edit services ipsec-vpn ipsec policy <i>policy-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define a list of proposals to include in the IKE or IPsec policy. |
| Options | <i>proposal-names</i> —List of IKE or IPsec proposal names. |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring IKE Proposals on page 405• Configuring IPsec Proposals on page 415 |

protocol (Applications)

| | |
|----------------------------|---|
| Syntax | <code>protocol type;</code> |
| Hierarchy Level | [edit applications application application-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Networking protocol type or number. |
| Options | <p>type—Networking protocol type. The following text values are supported:</p> <ul style="list-style-type: none"> ah egp esp gre icmp icmp6 igmp ipip ospf pim rsvp tcp udp |



NOTE: IP version 6 (IPv6) is not supported as a network protocol in application definitions.

| | |
|---------------------------------|--|
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • ALG Descriptions on page 277 • Configuring Application Sets on page 303 • Configuring Application Protocol Properties on page 303 • Examples: Configuring Application Protocols on page 321 • Verifying the Output of ALG Sessions |

protocol (IPSec)

| | |
|---------------------------------|--|
| Syntax | <code>protocol (ah esp bundle);</code> |
| Hierarchy Level | [edit services ipsec-vpn ipsec proposal <i>proposal-name</i>],
[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then manual direction <i>direction</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define an IPsec protocol for a dynamic or manual SA. |
| Options | ah —Authentication Header protocol.

esp —Encapsulating Security Payload protocol.

bundle —AH and ESP protocol. |
| Required Privilege Level | admin —To view this statement in the configuration.
admin-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Security Associations on page 385 |

ptsp-rules

| | |
|---------------------------------|--|
| Syntax | <code>(ptsp-rules <i>rule-name</i> ptsp-rules-sets <i>rule-set-name</i>);</code> |
| Hierarchy Level | [edit services service-set <i>service-set-name</i>] |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Specify the PTSP rules or rule set included in this service set. You can configure multiple rules but only one rule set for each service. |
| Options | <i>rule-name</i> —Identifier for the collection of terms that constitute this rule.

<i>rule-set-name</i> —Identifier for the set of rules to be included. |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Sets to be Applied to Services Interfaces on page 31 |

queues

| | |
|---------------------------------|--|
| Syntax | <code>queues [<i>queue-numbers</i>];</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression <i>rtp</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression <i>rtp</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For voice services interfaces only, assign queue numbers on which RTP compression takes place. |
| Options | queues <i>queue-numbers</i> —Assign one or more of the following queues: q0 , q1 , q2 , and q3 . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Compression of Voice Traffic on page 623 |

receive-window

| | |
|----------------------------|--|
| Syntax | <code>receive-window <i>packets</i>;</code> |
| Hierarchy Level | [edit services l2tp tunnel-group <i>name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the size of the receive window for L2TP tunnels, which limits the number of packets the server processes concurrently. |



NOTE: This statement is not supported for L2TP LNS on MX Series routers.

| | |
|---------------------------------|--|
| Options | packets —Maximum number of packets the receive window can hold at one time.
Default: 16 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Window Size for L2TP Tunnels on page 643 |

redistribute-all-traffic (Aggregated Multiservices)

| | |
|---------------------------------|---|
| Syntax | <code>redistribute-all-traffic {
 enable-rejoin;
}</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> load-balancing-options member-failure-options] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | <p>Enable the option to redistribute traffic of a failed active member to the other active members.</p> <p>For many-to-one (N:1) high availability support for Network Address Translation (NAT), the traffic for the failed member is automatically redistributed to the other active members.</p> <p>The remaining statement is explained separately.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Understanding Aggregated Multiservices Interfaces on page 599• Example: Configuring an Aggregated Multiservices Interface (AMS) on page 608• member-failure-options (Aggregated Multiservices) on page 1408 |

redundancy-options (Adaptive Services Interfaces)

| | |
|---------------------------------|---|
| Syntax | <pre> redundancy-options { primary sp-fpc/pic/port; secondary sp-fpc/pic/port; hot-standby } </pre> |
| Hierarchy Level | [edit interfaces rspnumber]
[edit interfaces rmsnumber] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the primary and secondary (backup) adaptive services interfaces. |
| Options | The remaining statements are explained separately. |
| Usage Guidelines | See “Configuring AS or Multiservices PIC Redundancy” on page 41 . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring AS or Multiservices PIC Redundancy on page 41 |

redundancy-options (Link Services IQ PIC Interfaces)

| | |
|---------------------------------|--|
| Syntax | <pre> redundancy-options { (hot-standby warm-standby); primary lsq-fpc/pic/port; secondary lsq-fpc/pic/port; } </pre> |
| Hierarchy Level | [edit interfaces rlsqnumber] |
| Release Information | Statement introduced in Junos OS Release 7.6. |
| Description | Specify the primary and secondary (backup) Link Services IQ PIC interfaces. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 548 |

(reflexive | reverse)

| | |
|---------------------------------|--|
| Syntax | <pre>(reflexive reverse) {
 application-profile profile-name;
 dscp (alias bits);
 forwarding-class class-name;
 syslog;
}</pre> |
| Hierarchy Level | [edit services cos rule rule-name term term-name then] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | <p>reflexive—Applies the equivalent opposing CoS action to flows in the opposite direction.</p> <p>reverse—Allows you to define CoS behavior for flows in the reverse direction.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring CoS Rules• Configuring Reflexive and Reverse CoS Rule Actions on page 518 |

rejoin-timeout (Aggregated Multiservices)

| | |
|---------------------------------|--|
| Syntax | <code>rejoin-timeout <i>rejoin-timeout</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> load-balancing-options member-failure-options drop-member-traffic] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Configure the time by when failed members (members in the DISCARD state) should rejoin the aggregated Multiservices (AMS) interface automatically. All members that do not rejoin by the configured time are moved to the INACTIVE state and the traffic meant for each of the members is dropped. |
| Default | If you do not configure a value, the default value of 120 seconds is used. |
| Options | <i>rejoin-timeout</i> —Time, in seconds, by which a failed member must rejoin.
Default: 120 seconds |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Understanding Aggregated Multiservices Interfaces on page 599 • Example: Configuring an Aggregated Multiservices Interface (AMS) on page 608 • drop-member-traffic (Aggregated Multiservices) on page 1350 |

remote-gateway

| | |
|---------------------------------|---|
| Syntax | <code>remote-gateway <i>address</i>;</code> |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the remote address to which the IPsec traffic is directed. |
| Options | <i>address</i> —Remote IPv4 or IPv6 address. |
| Required Privilege Level | admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IPsec Rules on page 422 |

remote-id

| | |
|---------------------------------|--|
| Syntax | <pre>remote-id {
 any-remote-id;
 ipv4_addr [<i>values</i>];
 ipv6_addr [<i>values</i>];
 key_id [<i>values</i>];
}</pre> |
| Hierarchy Level | [edit services ipsec-vpn ikepolicy <i>policy-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
ipv6_addr option added in Junos OS Release 7.6.
any-remote-id option added in Junos OS Release 8.2. |
| Description | Define the remote identification values to which the IKE policy applies. |
| Options | any-remote-id —Allow any remote address to connect. This option is supported only in dynamic configurations and cannot be configured with specific values.

ipv4_addr [<i>values</i>] —Define one or more IPv4 address identification values.

ipv6_addr [<i>values</i>] —Define one or more IPv6 address identification values.

key_id [<i>values</i>] —Define one or more key identification values.

fqdn <i>fqdn</i> —Fully-qualified domain name. |
| Required Privilege Level | admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IKE Policies on page 409 |

request-url

| | |
|---------------------------------|--|
| Syntax | <pre>request-url <i>page-name</i> ;</pre> |
| Hierarchy Level | [edit services hcm url-rule <i>url-rule-name</i> term <i>term-num</i> from url <i>url_identifier</i>] |
| Release Information | Statement introduced in Junos OS Release 12.1. |
| Description | Specify a request-URL to match the term . A match for the term is considered when a URL matches any hostname and any request-URL within a term. |
| Options | <i>page-name</i> —Page name of the request URL. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

retransmit-interval (Services)

| | |
|----------------------------|---|
| Syntax | <code>retransmit-interval <i>seconds</i>;</code> |
| Hierarchy Level | [edit services l2tp tunnel-group <i>name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the maximum retransmit interval for L2TP tunnels. |



NOTE: This statement is not supported for L2TP LNS on MX Series routers.

| | |
|---------------------------------|--|
| Options | <p><i>seconds</i>—Interval, in seconds, after which the server retransmits data if no acknowledgment is received.</p> <p>Default: 30 seconds</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Timers for L2TP Tunnels on page 643 |

rpc-program-number

| | |
|---------------------------------|--|
| Syntax | <code>rpc-program-number <i>number</i>;</code> |
| Hierarchy Level | [edit applications application <i>application-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Remote procedure call (RPC) or Distributed Computing Environment (DCE) value. |
| Options | <p><i>number</i>—RPC or DCE program value.</p> <p>Range: 100,000 through 400,000</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • ALG Descriptions on page 277 • Configuring an RPC Program Number on page 320 • Examples: Configuring Application Protocols on page 321 • Verifying the Output of ALG Sessions |

rtp

| | |
|---------------------------------|--|
| Syntax | <pre>rtp {
 f-max-period <i>number</i>;
 maximum-contexts <i>number</i> <force>;
 port {
 minimum <i>port-number</i>;
 maximum <i>port-number</i>;
 }
 queues [<i>queue-numbers</i>];
}</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the RTP properties for voice services traffic.

The remaining statements are described separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Compression of Voice Traffic on page 623 |

rule (Services CoS)

```
Syntax  rule rule-name {
        match-direction (input | output | input-output);
        term term-name {
            from {
                application-sets set-name;
                applications [ application-names ];
                destination-address address;
                destination-prefix-list list-name <except>;
                source-address address;
                source-prefix-list list-name <except>;
            }
            then {
                application-profile profile-name;
                dscp (alias | bits);
                forwarding-class class-name;
                syslog;
                (reflexive | reverse) {
                    application-profile profile-name;
                    dscp (alias | bits);
                    forwarding-class class-name;
                    syslog;
                }
            }
        }
    }
```

Hierarchy Level [edit [services cos](#)],
[edit [services cos rule-set](#) *rule-set-name*]

Release Information Statement introduced in Junos OS Release 8.1.

Description Specify the rule the router uses when applying this service.

Options *rule-name*—Identifier for the collection of terms that constitute this rule.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring CoS Rules on page 514](#)

rule (Services IDS)

```

Syntax  rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
        from {
            application-sets set-name;
            applications [ application-names ];
            destination-address (address | any-unicast) <except>;
            destination-address-range low minimum-value high maximum-value <except>;
            source-address (address | any-unicast) <except>;
            source-address-range low minimum-value high maximum-value <except>;
        }
        then {
            aggregation {
                destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
                source-prefix prefix-value | source-prefix-ipv6 prefix-value;
            }
            (force-entry | ignore-entry);
            logging {
                syslog;
                threshold rate;
            }
            session-limit {
                by-destination {
                    hold-time seconds;
                    maximum number;
                    packets number;
                    rate number;
                }
                by-pair {
                    hold-time seconds;
                    maximum number;
                    packets number;
                    rate number;
                }
                by-source {
                    hold-time seconds;
                    maximum number;
                    packets number;
                    rate number;
                }
            }
            syn-cookie {
                mss value;
                threshold rate;
            }
        }
    }
}

```

Hierarchy Level [edit [services](#) *ids*],
 [edit [services](#) *ids* [rule-set](#) *rule-set-name*]

| | |
|------------------------------|---|
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the rule the router uses when applying this service. |
| Options | <i>rule-name</i> —Identifier for the collection of terms that constitute this rule. |
| Required Privilege | interface—To view this statement in the configuration. |
| Level | interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IDS Rules on page 355 |

rule (Services IPsec VPN)

```
Syntax  rule rule-name {
        match-direction (input | output);
        term term-name {
            from {
                destination-address address;
                ipsec-inside-interface interface-name;
                source-address address;
            }
            then {
                anti-replay-window-size bits;
                backup-remote-gateway address;
                clear-dont-fragment-bit;
                dynamic {
                    ike-policy policy-name;
                    ipsec-policy policy-name;
                }
                initiate-dead-peer-detection;
                manual {
                    direction (inbound | outbound | bidirectional) {
                        authentication {
                            algorithm (hmac-md5-96 | hmac-sha1-96);
                            key (ascii-text key | hexadecimal key);
                        }
                        auxiliary-spi spi-value;
                        encryption {
                            algorithm algorithm;
                            key (ascii-text key | hexadecimal key);
                        }
                        protocol (ah | bundle | esp);
                        spi spi-value;
                    }
                }
                no-anti-replay;
                remote-gateway address;
                syslog;
                tunnel-mtu bytes;
            }
        }
    }
```

Hierarchy Level [edit [services](#) ipsec-vpn],
[edit [services](#) ipsec-vpn [rule-set](#) *rule-set-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the rule the router uses when applying this service.

Options *rule-name*—Identifier for the collection of terms that comprise this rule.

The remaining statements are explained separately.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring IPsec Rules on page 422](#)
- [Configuring IPsec Rule Sets on page 429](#)
- [Configuring Security Associations on page 385](#)

rule (Services NAT)

```
Syntax  rule rule-name {
        match-direction (input | output);
        term term-name {
            from {
                application-sets set-name;
                applications [ application-names ];
                destination-address (address | any-unicast) <except>;
                destination-address-range low minimum-value high maximum-value <except>;
                source-address (address | any-unicast) <except>;
                source-address-range low minimum-value high maximum-value <except>;
            }
            then {
                no-translation;
                translated {
                    address-pooling paired;
                    destination-pool nat-pool-name;
                    destination-prefix destination-prefix; destination-prefix;
                    dns-alg-pool dns-alg-pool;
                    dns-alg-prefix dns-alg-prefix;
                    filtering-type endpoint-independent;
                    mapping-type endpoint-independent;
                    overload-pool overload-pool;
                    overload-prefix overload-prefix;
                    source-pool nat-pool-name;
                    source-prefix source-prefix;
                    translation-type (basic-nat-pt | basic-nat44 | basic-nat66 | dnat-44 | dynamic-nat44
                        | napt-44 | napt-66 | napt-pt | stateful-nat64 | twice-basic-nat-44 |
                        twice-dynamic-nat-44 | twice-napt-44);
                }
            }
            syslog;
        }
    }
```

Hierarchy Level [edit [services](#) nat],
[edit [services](#) nat [rule-set](#) rule-set-name]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the rule the router uses when applying this service.



NOTE: You are limited to a maximum of 200 terms for a NAT rule that is applied to an inline services (type si) interface. If you specify more than 200 terms, you will receive following error when you commit the configuration:

```
[edit]
' service-set service-set-name'
  NAT rule rule-name with more than 200 terms is disallowed for
  si-n/n/n.n
error: configuration check-out failed
```


Options *rule-name*—Identifier for the collection of terms that make up this rule.

The remaining statements are explained separately.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation • [Network Address Translation Rules Overview on page 69](#)

rule (Services Stateful Firewall)

Syntax

```
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application-sets set-name;
      applications [ application-names ];
      destination-address (address | any-unicast) <except>;
      destination-address-range low minimum-value high maximum-value <except>;
      destination-prefix-list list-name <except>;
      source-address (address | any-unicast) <except>;
      source-address-range low minimum-value high maximum-value <except>;
      source-prefix-list list-name <except>;
    }
    then {
      (accept | discard | reject);
      syslog;
    }
  }
}
```

Hierarchy Level [edit [services](#) stateful-firewall],
 [edit [services](#) stateful-firewall [rule-set](#) rule-set-name]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the rule the router uses when applying this service.

Options *rule-name*—Identifier for the collection of terms that constitute this rule.

The remaining statements are explained separately.

Usage Guidelines See “[Configuring Stateful Firewall Rules](#)” on page 331.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

rule (Softwire)

| | |
|---------------------------------|---|
| Syntax | <pre>rule <i>rule-name</i> {
 match-direction (input output);
 term <i>term-name</i> {
 then {
 (ds-lite <i>ds-lite-softwire-concentrator</i> v6rd <i>v6rd-softwire-concentrator</i>);
 }
 }
}</pre> |
| Hierarchy Level | [edit services softwire],
[edit services softwire rule-set <i>rule-set-name</i>] |
| Release Information | Statement introduced in Junos OS Release 10.4. |
| Description | Configure a rule to apply a softwire concentrator for a flow. |
| Options | <p><i>rule-name</i>—Identifier for the collection of terms that constitute this rule.</p> <p><i>input</i>—Apply the rule match on the input side of the interface.</p> <p><i>output</i>—Apply the rule match on the output side of the interface.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Softwire Rules on page 219 |

rule-set (Services CoS)

| | |
|---------------------------------|---|
| Syntax | <pre>rule-set <i>rule-set-name</i> {
 [rule <i>rule-name</i>];
}</pre> |
| Hierarchy Level | [edit services cos] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | Specify the rule set the router uses when applying this service. |
| Options | <i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring CoS Rule Sets |

rule-set (Services IDS)

| | |
|---------------------------------|---|
| Syntax | <code>rule-set <i>rule-set-name</i> {
 [<i>rule</i> <i>rule-names</i>];
}</code> |
| Hierarchy Level | [edit services ids] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the rule set the router uses when applying this service. |
| Options | <i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IDS Rule Sets on page 363 |

rule-set (Services IPsec VPN)

| | |
|---------------------------------|---|
| Syntax | <code>rule-set <i>rule-set-name</i> {
 [<i>rule</i> <i>rule-names</i>];
}</code> |
| Hierarchy Level | [edit services ipsec-vpn] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the rule set the router uses when applying this service. |
| Options | <i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring IPsec Rules on page 422 |

rule-set (Services NAT)

| | |
|---------------------------------|---|
| Syntax | <code>rule-set <i>rule-set-name</i> {
 [<i>rule</i> <i>rule-names</i>];
}</code> |
| Hierarchy Level | [edit services nat] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the rule set the router uses when applying this service. |
| Options | <i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |

rule-set (Services Stateful Firewall)

| | |
|---------------------------------|---|
| Syntax | <code>rule-set <i>rule-set-name</i> {
 [<i>rule</i> <i>rule-names</i>];
}</code> |
| Hierarchy Level | [edit services stateful-firewall] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the rule set the router uses when applying this service. |
| Options | <i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set. |
| Usage Guidelines | See “ Configuring Stateful Firewall Rule Sets ” on page 335. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

rule-set (Softwire)

| | |
|---------------------------------|---|
| Syntax | <code>rule-set <i>rule-set-name</i> {
 rule <i>rule-name</i>;
}</code> |
| Hierarchy Level | [edit services softwire] |
| Release Information | Statement introduced in Junos OS Release 10.4. |
| Description | Specify the rule set the router uses when applying this service. |
| Options | <i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Softwire Rules on page 219 |

secondary (Adaptive Services Interfaces)

| | |
|---------------------------------|---|
| Syntax | <code>secondary <i>interface-name</i>;</code> |
| Hierarchy Level | [edit interfaces (rsp0 rsp1) redundancy-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the secondary (backup) adaptive services interface. |
| Options | <i>interface-name</i> —The identifier for the adaptive services interface, which must be of the form <i>sp-fpc/pic/port</i> . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring AS or Multiservices PIC Redundancy on page 41 |

secondary (Link Services IQ PIC Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>secondary <i>interface-name</i>;</code> |
| Hierarchy Level | [edit interfaces <i>rlsnumber</i> redundancy-options] |
| Release Information | Statement introduced in Junos OS Release 7.6. |
| Description | Specify the secondary (backup) Link Services IQ PIC interface. |
| Options | <i>interface-name</i> —The identifier for the Link Services IQ PIC interface, which must be of the form <i>lsq-fpc/pic/port</i> . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 548 |

secure-nat-mapping

| | |
|---------------------------------|--|
| Syntax | <pre>secure-nat-mapping {
 mapping-refresh (inbound outbound inbound-outbound);
 eif-flow-limit <i>number-of-flows</i>
}</pre> |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated] |
| Release Information | Statement introduced in Junos OS Release 12.3 |
| Description | Specify configuration options that help prevent or minimize the effect of attempted denial of service (DOS) attacks for NAT operations. |
| Options | The statements are explained separately.

— |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Protecting CGN Devices Against Denial of Service (DOS) Attacks on page 241 |

secured-port-block-allocation

Syntax `secured-port-block-allocation {
 active-block-timeout timeout-seconds;
 block-size block-size;
 max-blocks-per-address max-blocks;
 }`

Hierarchy Level [edit [services](#) nat [pool](#) *pool-name* port]

Release Information Statement introduced in Junos OS Release 11.2.

Description When you use block allocation, one or more blocks of ports in a NAT pool address range are available for assignment to a subscriber.



NOTE: If you define the session lifetime globally for a Multiservices (ms-) interface (by using the `session-timeout seconds` statement at the [edit `interfaces interface-name services-options`] hierarchy level), the session is terminated even if traffic continues to flow beyond that time period. When continuous traffic transmission occurs, the session is reset immediately after the timeout period. When the session timeout value is the same as the timeout value for active port block allocation, it might be possible that the system does not determine that the active port block timeout period has elapsed. As a result, for the first allocation of a port block after the active block timeout occurs, the same block that was previously used might be used for allocation. However, for the subsequent allocation of a port block, the system identifies the active block timeout value correctly and allocates a port from a new block. This behavior is expected when the session timeout and port block timeout values are identical. To avoid this problem, we recommend that you configure different values for session timeout and port block timeout so that the `JSERVICES_NAT_PORT_BLOCK_ALLOC` system logging message is generated at correct intervals of the active port block timeout value.

Options *block-size*—Number of ports included in a block.

Default: 128

Range: 1 through 32,000

max-blocks—Maximum number of blocks that can be allocated to a user address.

Default: 8

Range: 1 to 512

timeout-seconds—Interval, in seconds, during which a block is active. After timeout, a new block is allocated, even if ports are available in the active block.

Default: 120

Range: 0 through 86400. When you specify 0, the active block transitions to inactive only when it runs out of ports and a new block is allocated. Any inactive block without any ports in use will be freed to the NAT pool.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation • [Configuring Address Pools for Network Address Port Translation \(NAPT\) Overview on page 117](#)

server (pcp)

Syntax `server server-name {
 ipv4-address ipv4-address;
 ipv6-address ipv6-address;
 software-concentrator software-concentrator-name;
 mapping-lifetime-min mapping-lifetime-min;
 mapping-lifetime-max mapping-lifetime-max;
 short-lifetime-error short-lifetime-error;
 long-lifetime-error long-lifetime-error;
 nat-options {
 pool pool-name ;
 }
 pcp-options {
 third-party
 prefer-failure
 }
 max-mapping-per-client max-mapping-per-client;
}`

Hierarchy Level [edit services pcp]

Release Information Statement introduced in Junos OS Release 13.2R1.

Description Configure PCP server options.

Options *ipv4-address*—IPv4 address of the PCP server.

ipv6-address—IPv6 address of the PCP server.

software-concentrator-name—Softwire concentrator name whose softwire-address is used in creating PCP mappings. The PCP server address must be the same as the softwire-concentrator address.

mapping-lifetime-min—Minimum lifetime, in seconds, for PCP mapping. If a PCP client requests a lifetime less than the minimum configured, the server will assign a minimum lifetime and respond accordingly.

Default: 300 seconds

Range: 120 through 3600 seconds

mapping-lifetime-max mapping-lifetime-max—Maximum lifetime, in seconds, for PCP mapping. If the PCP client requests a lifetime less than the maximum configured, the server will assign the maximum lifetime and respond accordingly.

Default: 86,400 seconds

Range: 3600 through 2147483647 seconds

short-lifetime-error short-lifetime-error—Certain error opcodes mentioned in section 2 are classified as short lifetime errors. In case of these errors, the PCP server will use the value configured with this option to respond to the PCP client.

Default: 30 seconds

Range: 15 through 300 seconds

long-lifetime-error—Certain error opcodes mentioned in section 2 are classified as long lifetime errors. In case of these errors, the PCP server will use the value configured with this option to respond to the PCP client.

Default: 1800 seconds

Range: 900 through 18,000 seconds

max-mapping-per-client *number-of-mappings*—Maximum number of PCP mappings that the PCP client can request.

Default: 32

Range: 1 through 32

The other statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Port Control Protocol on page 167](#)

service

Syntax

```
service {  
  input {  
    [ service-set service-set-name <service-filter filter-name> ];  
    post-service-filter filter-name;  
  }  
  output {  
    [ service-set service-set-name <service-filter filter-name> ];  
  }  
}
```

Hierarchy Level [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** inet],
[edit logical-systems *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family** inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the service sets and filters to be applied to an interface.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Applying Filters and Services to Interfaces on page 38](#)

service-domain

| | |
|---------------------------------|---|
| Syntax | <code>service-domain (inside outside);</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the service interface domain. If you specify this interface using the next-hop-service statement at the [edit services service-set <i>service-set-name</i>] hierarchy level, the interface domain must match that specified with the inside-service-interface and outside-service-interface statements. |
| Options | inside —Interface used within the network.

outside —Interface used outside the network. |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Address and Domain for Services Interfaces on page 45 |


service-filter (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>service-filter <i>filter-name</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service (input output) service-set <i>service-set-name</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service (input output) service-set <i>service-set-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the filter to be applied to traffic before it is accepted for service processing. Configuration of a service filter is optional; if you include the service-set statement without a service-filter definition, Junos OS assumes the match condition is true and selects the service set for processing automatically. |
| Options | <i>filter-name</i> —Identifies the filter to be applied in service processing. |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Applying Filters and Services to Interfaces on page 38 • Junos OS Services Interfaces Library for Routing Devices |

service-interface (Adaptive Services Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>service-interface <i>interface-name</i>;</code> |
| Hierarchy Level | [edit services service-set <i>service-set-name</i> interface-service] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the name for the adaptive services interface associated with an interface-wide service set. |
| Options | interface-name —Identifier of the service interface. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Sets to be Applied to Services Interfaces on page 31 |

service-interface (L2TP Processing)

| | |
|--|---|
| Syntax | <code>service-interface <i>interface-name</i>;</code> |
| Hierarchy Level | [edit services l2tp tunnel-group name] |
| Release Information | Statement introduced before Junos OS Release 7.4.
<i>si-fpc/pic/port</i> option added in Junos OS Release 11.4. |
| Description | Specify the service interface responsible for handling L2TP processing. |
| <div> NOTE: On MX Series routers, the service interface configuration is required for static LNS sessions. Either the service interface configuration or the service device pool configuration can be used for dynamic LNS sessions.</div> | |
| Options | interface-name —Name of the service interface. The interface type depends on the line card as follows: <ul style="list-style-type: none">• <i>sp-fpc/pic/port</i>—On AS or Multiservices PICs on M7i, M10i, and M120 routers.• <i>si-fpc/pic/port</i>—On MPCs on MX Series routers. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Local Gateway Address and PIC on page 642• Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces |

service-set (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>service-set <i>service-set-name</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service (input output)],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service (input output)] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define one or more service sets to be applied to an interface. If you define multiple service sets, the router software evaluates the filters in the order in which they appear in the configuration. |
| Options | <i>service-set-name</i> —Identifies the service set. |
| Required Privilege Level | System—To view this statement in the configuration.
System-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Applying Filters and Services to Interfaces on page 38 |

service-set (Services)

Syntax `service-set service-set-name {`
 `allow-multicast;`
 `extension-service service-name {`
 `provider-specific-rules-configuration;`
 `}`
 `(ids-rules rule-name | ids-rule-sets rule-set-name);`
 `interface-service {`
 `service-interface interface-name;`
 `}`
 `ipsec-vpn-options {`
 `anti-replay-window-size bits;`
 `clear-dont-fragment-bit;`
 `ike-access-profile profile-name;`
 `local-gateway address;`
 `no-anti-replay;`
 `passive-mode-tunneling;`
 `trusted-ca [ca-profile-names];`
 `tunnel-mtu bytes;`
 `}`
 `ip-reassembly-rules rule-name;`
 `(ipsec-vpn-rules rule-name | ipsec-vpn-rule-sets rule-set-name);`
 `max-flows number;`
 `max-drop-flows {`
 `ingress ingress-flows;`
 `egress egress-flows;`
 `}`
 `nat-options {`
 `land-attack-check (ip-only | ip-port);`
 `max-sessions-per-subscriber session-number;`
 `stateful-nat64{`
 `clear-dont-fragment-bit;`
 `}`
 `}`
 `(nat-rules rule-name | nat-rule-sets rule-set-name);`
 `next-hop-service {`
 `inside-service-interface interface-name.unit-number;`
 `outside-service-interface interface-name.unit-number;`
 `outside-service-interface-type local;`
 `service-interface-pool name;`
 `}`
 `(pgcp-rules rule-name | pgcp-rule-sets rule-set-name);`
 `(ptsp-rules rule-name | ptsp-rule-sets rule-set-name);`
 `service-set-options {`
 `bypass-traffic-on-exceeding-flow-limits;`
 `bypass-traffic-on-pic-failure;`
 `enable-asymmetric-traffic-processing;`
 `routing-engine-services;`
 `support-uni-directional-traffic;`
 `}`
 `snmp-trap-thresholds{`
 `flows high high-threshold | low low-threshold;`
 `nat-address-port high-threshold | low low-threshold;`

```

    }
  }
  software-options {
    dslite-ipv6-prefix-length dslite-ipv6-prefix-length;
  }
  (software-rules rule-name | software-rule-sets rule-set-name);
  (stateful-firewall-rules rule-name | stateful-firewall-rule-sets rule-set-name);
  syslog {
    host hostname {
      class {
        alg-logs;
        ids-logs;
        nat-logs;
        packet-logs;
        pcp-logs;
        session-logs <open | close>;
        stateful-firewall-logs ;
      }
      services severity-level;
      facility-override facility-name;
      interface-service prefix-value;
    }
  }
}

```

| | |
|--------------------------|---|
| Hierarchy Level | [edit services] |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>pgcp-rules and pgcp-rule-sets options added in Junos OS Release 8.4.</p> <p>server-set-options option added in Junos OS Release 10.1.</p> <p>ptsp-rules and ptsp-rule-sets options added in Junos OS Release 10.2.</p> <p>software-rules and clear-rule-sets options added in Junos OS Release 10.4.</p> <p>software-options option added in Junos OS Release 14.1.</p> |
| Description | Define the service set. |
| Options | <p><i>service-set-name</i>—Name of the service set.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> <i>Service Set Properties</i> |

service-set-options

| | |
|---------------------------------|---|
| Syntax | <pre>service-set-options {
 bypass-traffic-on-exceeding-flow-limits;
 bypass-traffic-on-pic-failure;
 enable-asymmetric-traffic-processing;
 header-integrity-check
 routing-engine-services;
 support-uni-directional-traffic;
}</pre> |
| Hierarchy Level | [edit services service-set] |
| Release Information | Statement introduced in Junos OS Release 10.1. The enable-asymmetric-traffic-processing and the support-uni-directional-traffic options were added in Junos OS Release 11.2. The routing-engine-services option was added in Junos OS Release 15.1. |
| Description | Specify the service set options to apply to a service set. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Sets to be Applied to Services Interfaces on page 31• Configuring APPID Support for Unidirectional Traffic on page 685 |

session-limit

Syntax

```
session-limit {
  by-destination {
    hold-time seconds;
    maximum number;
    packets number;
    rate number;
  }
  by-pair {
    hold-time seconds;
    maximum number;
    packets number;
    rate number;
  }
  by-source {
    hold-time seconds;
    maximum number;
    packets number;
    rate number;
  }
}
```

Hierarchy Level [edit [services](#) ids [rule](#) *rule-name* [term](#) *term-name* [then](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable flow limitation by configuring thresholds on source, destination, or stateful firewall and network address translation (NAT) paired traffic flows.

Options The remaining statements are described separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Actions in IDS Rules on page 358](#)

set-dont-fragment-bit (Services Set)

| | |
|---------------------------------|---|
| Syntax | set-dont-fragment-bit; |
| Hierarchy Level | [edit services service-set <i>service-set-name</i> ipsec-vpn-options] |
| Release Information | Statement introduced in Junos OS Release 14.1. |
| Description | Configure the do not fragment (DF) bit in only the outer header of the IPsec packet and leave the inner header unmodified for dynamic endpoint tunnels. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. These settings apply for dynamic endpoint tunnels and not for static tunnels, for which you need to include the set-dont-fragment-bit statement at the [edit services ipsec-vpn rule rule-name term term-name then] hierarchy level to set the DF bit in the outer header of the IPv4 packets that enter the static IPsec tunnel. This functionality is supported on MX Series routers with MS-MICs and MS-MPCs. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Service Sets on page 430• Configuring IPsec Rules |

set-dont-fragment-bit (Services IPsec VPN)

| | |
|---------------------------------|---|
| Syntax | set-dont-fragment-bit; |
| Hierarchy Level | [edit services ipsec-vpn rule rule-name term term-name then] |
| Release Information | Statement introduced in Junos OS Release 14.1. |
| Description | Configure the do not fragment (DF) bit in only the outer header of the IPsec packet and leave the inner header unmodified. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. These settings apply for static endpoint tunnels and not for dynamic tunnels, for which you need to include the set-dont-fragment-bit statement at the [edit services service-set service-set-name ipsec-vpn-options] hierarchy level to set the DF bit in the outer header of the IPv4 packets that enter the dynamic IPsec tunnel. This functionality is supported on MX Series routers with MS-MICs and MS-MPCs. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Rules on page 422 |

sip-call-hold-timeout

| | |
|---------------------------------|---|
| Syntax | <code>sip-call-hold-timeout seconds;</code> |
| Hierarchy Level | [edit applications application application-name] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Timeout period for SIP calls placed on hold, in seconds. |
| Options | seconds —Length of time the application holds a SIP call open before it times out.
Default: 7200 seconds
Range: 0 through 36,000 seconds (10 hours) |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• ALG Descriptions on page 277• Configuring SIP on page 303• Examples: Configuring Application Protocols on page 321• Verifying the Output of ALG Sessions |

sip

| | |
|---------------------------------|--|
| Syntax | <pre>sip {
 video {
 dscp (alias bits);
 forwarding-class class-name;
 }
 voice {
 dscp (alias bits);
 forwarding-class class-name;
 }
}</pre> |
| Hierarchy Level | [edit services cos application-profile <i>profile-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | Set the appropriate dscp and forwarding-class value for SIP traffic. |
| Default | By default, the system will not alter the DSCP or forwarding class for SIP traffic.

The remaining statements are explained separately. |
| Usage Guidelines | See “ Configuring CoS Rules ” on page 514 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

snmp-command

| | |
|---------------------------------|---|
| Syntax | <pre>snmp-command <i>command</i>;</pre> |
| Hierarchy Level | [edit applications application <i>application-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | SNMP command format. |
| Options | command —Supported commands are SNMP get , get-next , set , and trap . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• ALG Descriptions on page 277• Configuring an SNMP Command for Packet Matching on page 320• Examples: Configuring Application Protocols on page 321• Verifying the Output of ALG Sessions |

software-concentrator

Syntax

```
software-concentrator {
  ds-lite ds-lite-software-concentrator {
    auto-update-mtu;
    flow-limit flow-limit | session-limit-per-prefix session-limit-per-prefix;
    mtu-v6 mtu-v6;
    software-address address;
  }
  v6rd v6rd-software-concentrator {
    ipv4-prefix ipv4-prefix;
    v6rd-prefix ipv6-prefix;
    mtu-v4 mtu-v4;
  }
}
```

Hierarchy Level [edit services software]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure settings for a software concentrator.

Options The remaining statements are explained separately.

Required Privilege Level

| |
|---|
| interface—To view this statement in the configuration. |
| interface-control—To add this statement to the configuration. |

Related Documentation

- [Configuring a DS-Lite Software Concentrator on page 227](#)
- [Configuring a 6rd Software Concentrator on page 245](#)


software-options

| | |
|---------------------------------|---|
| Syntax | <code>software-options {
 dslite-ipv6-prefix-length <i>dslite-ipv6-prefix-length</i> ;
}</code> |
| Hierarchy Level | [edit services service-set <i>service-set-name</i>] |
| Release Information | Statement introduced in Junos OS Release 14.1. |
| Description | Specify the IPv6 prefix length associated with a subscriber's basic broadband bridging device that is subject to a limited number of sessions. |
| Options | <i>dslite-ipv6-prefix-length</i> —Subnet prefix representing the size of the subnet subject to session limitation.
Values: 56, 64, 96, 128
Default: 0—no limitation. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• DS-Lite Per Subnet Limitation Overview on page 242 |

software-rules

| | |
|---------------------------------|---|
| Syntax | <code>(software-rule <i>rule-name</i> software-rule-sets <i>rule-set-name</i>);</code> |
| Hierarchy Level | [edit services service-set <i>service-set-name</i>] |
| Release Information | Statement introduced in Junos OS Release 10.4. |
| Description | Specify the DS-Lite or 6rd rules or rule set included in this service set. You can configure multiple rules; however, you can only configure one rule set for each service set. |
| Options | <i>rule-name</i> —Identifier for the collection of terms that constitute this rule.
<i>rule-set-name</i> —Identifier for the set of rules to be included. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Rules on page 36 |

source-address (Service Sets)

| | |
|---|---|
| Syntax | <code>source-address source-address</code> |
| Hierarchy Level | [edit services service-set service-set-name syslog host hostname] |
| Release Information | Statement introduced in Junos OS Release 13.1. |
| Description | Specify a source address to record in system log messages that are directed to a remote machine specified in the hostname statement. |
| <div>  <p>NOTE: The supported interfaces are ms, rms, and mams interfaces. If you do not specify the interface parameter, the command loops on all supported interfaces.</p> </div> | |
| Options | source-address —A valid IP address, which is recorded as the message source in messages sent to the remote machines specified in the host hostname statement |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring System Logging for Service Sets on page 47 • host on page 1378 • service-set on page 1277 |

source-address (Services CoS)

| | |
|--------------------------|---|
| Syntax | <code>source-address address;</code> |
| Hierarchy Level | [edit services cos rule rule-name term term-name from] |
| Release Information | Statement introduced in Junos OS Release 8.1.
address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5. |
| Description | Source address for rule matching. |
| Options | address —Source IPv4 or IPv6 address or prefix value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Match Conditions in a CoS Rule • Configuring Match Conditions In CoS Rules on page 515 |

source-address (Services IDS)

| | |
|---------------------------------|---|
| Syntax | <code>source-address (<i>address</i> any-unicast) <except>;</code> |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4.
<i>address</i> option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the source address for rule matching. |
| Options | <i>address</i> —Source IPv4 or IPv6 address or prefix value.
<i>any-unicast</i> —Any unicast packet.
<i>except</i> —(Optional) Exempt the specified address, prefix, or unicast packets from rule matching. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Match Conditions in IDS Rules on page 357 |

source-address (Services IPsec VPN)

| | |
|---------------------------------|---|
| Syntax | <code>source-address <i>address</i>;</code> |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the source address for rule matching. |
| Options | <i>address</i> —Source IP address. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Rules on page 422 |

source-address (Services NAT)

| | |
|---------------------------------|--|
| Syntax | source-address (<i>address</i> any-unicast) <except>; |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4.
any-unicast and except options introduced in Junos OS Release 7.6.
address option enhanced to support IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the source address for rule matching. |
| Options | address —Source IPv4 or IPv6 address or prefix value.

any-unicast —Any unicast packet.

except —(Optional) Prevent the specified address or unicast packets from being translated. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Network Address Translation Rules Overview on page 69 |

source-address (Services Stateful Firewall)

| | |
|---------------------------------|---|
| Syntax | source-address (<i>address</i> any-unicast) <except>; |
| Hierarchy Level | [edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced before Junos OS Release 7.4.
any-unicast and except options introduced in Junos OS Release 7.6.
address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5. |
| Description | Source address for rule matching. |
| Options | address —Source IPv4 or IPv6 address or prefix value.

any-unicast —Any unicast packet.

except —(Optional) Exclude the specified address, prefix, or unicast packets from rule matching. |
| Usage Guidelines | See “ Configuring Match Conditions in Stateful Firewall Rules ” on page 332. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

source-address-range (Services IDS)

| | |
|--------------------------|--|
| Syntax | source-address-range low <i>minimum-value</i> high <i>maximum-value</i> <except>; |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 7.6.
<i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the source address range for rule matching. |
| Options | <i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range.
<i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range.
<i>except</i> —(Optional) Exempt the specified address range from rule matching. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Match Conditions in IDS Rules on page 357 |

source-address-range (Services NAT)

| | |
|--------------------------|--|
| Syntax | source-address-range low <i>minimum-value</i> high <i>maximum-value</i> <except>; |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 7.6.
<i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the source address range for rule matching. |
| Options | <i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range.
<i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range.
<i>except</i> —(Optional) Prevent the specified address range from being translated. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |

source-address-range (Services Stateful Firewall)

| | |
|---------------------------------|---|
| Syntax | source-address-range low <i>minimum-value</i> high <i>maximum-value</i> <except>; |
| Hierarchy Level | [edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 7.6.
<i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5. |
| Description | Source address range for rule matching. |
| Options | <i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range.
<i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range.
<i>except</i> —(Optional) Exclude the specified address, prefix, or unicast packets from rule matching. |
| Usage Guidelines | See “ Configuring Match Conditions in Stateful Firewall Rules ” on page 332. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

source-pool

| | |
|---------------------------------|--|
| Syntax | source-pool <i>nat-pool-name</i> ; |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the source address pool for translated traffic. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Network Address Translation Rules Overview on page 69 |

source-port

| | |
|---------------------------------|--|
| Syntax | source-port <i>port-number</i> ; |
| Hierarchy Level | [edit applications application application-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Source port identifier. |
| Options | <i>port-value</i> —Identifier for the port. For a complete list, see “ Configuring Source and Destination Ports ” on page 309. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• ALG Descriptions on page 277• Configuring Application Protocol Properties on page 303• Configuring Source and Destination Ports on page 309• Verifying the Output of ALG Sessions |

source-prefix (Services IDS)

| | |
|---------------------------------|---|
| Syntax | source-prefix <i>prefix-value</i> ; |
| Hierarchy Level | [edit services ids rule rule-name term term-name then aggregation] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the prefix value for source IPv4 address aggregation. |
| Options | <i>prefix-value</i> —Integer value.
Range: 1 through 32 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Actions in IDS Rules on page 358 |

source-prefix (Services NAT)

| | |
|---------------------------------|--|
| Syntax | <code>source-prefix <i>source-prefix</i>;</code> |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated] |
| Release Information | Statement introduced in Junos OS Release 7.6.
<i>source-prefix</i> option enhanced to support IPv6 addresses in Junos OS Release 8.5. |
| Description | Specify the source prefix for translated traffic. |
| Options | <i>source-prefix</i> —IPv4 or IPv6 source prefix value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Network Address Translation Rules Overview on page 69 |

source-prefix-ipv6

| | |
|---------------------------------|---|
| Syntax | <code>source-prefix-ipv6 <i>prefix-value</i>;</code> |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> then aggregation] |
| Release Information | Statement introduced in Junos OS Release 8.5. |
| Description | Specify the prefix value for source IPv6 address aggregation. |
| Options | <i>prefix-value</i> —Integer value.
Range: 1 through 128 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Actions in IDS Rules on page 358 |

source-prefix-list (Services CoS)

| | |
|--------------------------|---|
| Syntax | source-prefix-list <i>list-name</i> <except>; |
| Hierarchy Level | [edit services cos rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 8.2. |
| Description | Specify the source prefix list for rule matching. You configure the prefix list by including the prefix-list statement at the [edit policy-options] hierarchy level. |
| Options | <i>list-name</i> —Destination prefix list.

except—(Optional) Exclude the specified prefix list from rule matching. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring CoS Rules on page 514• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> |

source-prefix-list (Services IDS)

| | |
|--------------------------|---|
| Syntax | source-prefix-list <i>list-name</i> <except>; |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 8.2. |
| Description | Specify the source prefix list for rule matching. You configure the prefix list by including the prefix-list statement at the [edit policy-options] hierarchy level. |
| Options | <i>list-name</i> —Destination prefix list.

except—(Optional) Exclude the specified prefix list from rule matching. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Match Conditions in IDS Rules on page 357• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> |


source-prefix-list (Services NAT)

| | |
|---------------------------------|--|
| Syntax | source-prefix-list <i>list-name</i> <except>; |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 8.2. |
| Description | Specify the source prefix list for rule matching. You configure the prefix list by including the prefix-list statement at the [edit policy-options] hierarchy level. |
| Options | <p>list-name—Destination prefix list.</p> <p>except—(Optional) Exclude the specified prefix list from rule matching.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Network Address Translation Rules Overview on page 69 • <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> |

source-prefix-list (Services Stateful Firewall)

| | |
|---------------------------------|---|
| Syntax | source-prefix-list <i>list-name</i> <except>; |
| Hierarchy Level | [edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 8.2. |
| Description | Specify the source prefix list for rule matching. You configure the prefix list by including the prefix-list statement at the [edit policy-options] hierarchy level. |
| Options | <p>list-name—Destination prefix list.</p> <p>except—(Optional) Exclude the specified prefix list from rule matching.</p> |
| Usage Guidelines | See “ Configuring Match Conditions in Stateful Firewall Rules ” on page 332. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> |

spi

| | |
|--|--|
| Syntax | <code>spi spi-value;</code> |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then manual direction <i>direction</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the SPI for an SA. |
| Options | spi-value —An arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet).
Range: 256 through 16,639 |
| <div> NOTE: Use the auxiliary SPI when you configure the protocol statement to use the bundle option.</div> | |
| Required Privilege Level | system—To view this statement in the configuration.
system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Security Associations on page 385 |

stateful-firewall-rules

| | |
|---------------------------------|---|
| Syntax | <code>(stateful-firewall-rules <i>rule-names</i> stateful-firewall-rule-sets <i>rule-set-name</i>);</code> |
| Hierarchy Level | [edit services service-set <i>service-set-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the stateful firewall rules or rule set included in this service set. You can configure multiple rules, but only one rule set for each service. |
| Options | rule-name —Identifier for the collection of terms that make up this rule.
rule-set-name —Identifier for the set of rules to be included. |
| Required Privilege Level | System—To view this statement in the configuration.
System-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Rules on page 36 |

syslog (Services CoS)

| | |
|---------------------------------|--|
| Syntax | syslog; |
| Hierarchy Level | [edit services cos rule <i>rule-name</i> term <i>term-name</i> then],
[edit services cos rule <i>rule-name</i> term <i>term-name</i> then (reflexive reverse)] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | Enable system logging. The system log information from the Adaptive Services or Multiservices PIC is passed to the kernel for logging in the /var/log directory. This setting overrides any syslog statement setting included in the service set or interface default configuration. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring Actions in a CoS Rule</i> • Configuring Actions in CoS Rules on page 516 |

syslog (Services IDS)

| | |
|---------------------------------|---|
| Syntax | syslog; |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> then logging] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Enable system logging. The system log information from the Adaptive Services or Multiservices Physical Interface Card (PIC) is passed to the kernel for logging in the /var/log directory. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Actions in IDS Rules on page 358 |

syslog (Services IPsec VPN)

| | |
|---------------------------------|--|
| Syntax | syslog; |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Enable system logging. The system log information for the Adaptive Services or Multiservices Physical Interface Card (PIC) is passed to the kernel for logging in the <code>/var/log</code> directory. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Rules on page 422 |

syslog (Services L2TP)

| | |
|----------------------------|--|
| Syntax | <pre>syslog {
 host <i>hostname</i> {
 services <i>severity-level</i>;
 facility-override <i>facility-name</i>;
 log-prefix <i>prefix-value</i>;
 }
}</pre> |
| Hierarchy Level | [edit services l2tp tunnel-group <i>group-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the generation of system log messages for L2TP services. System log information is passed to the kernel for logging in the <code>/var/log/l2tpd</code> directory. |



NOTE: This statement is not supported for L2TP LNS on MX Series routers.

| | |
|---------------------------------|--|
| Options | The remaining statements are described separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring System Logging of L2TP Tunnel Activity on page 644 |

syslog (Services NAT)

| | |
|---------------------------------|--|
| Syntax | syslog; |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Enable system logging. The system log information from the Multiservices PIC is passed to the kernel for logging in the /var/log directory. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Network Address Translation Rules Overview on page 69 |

syslog (Services Service Set)

| | |
|--------------------------|---|
| Syntax | <pre>syslog {
 host <i>hostname</i> {
 class {
 alg-logs;
 ids-logs;
 nat-logs;
 packet-logs;
 pcp-logs;
 session-logs <open close>;
 stateful-firewall-logs ;
 }
 services <i>severity-level</i>;
 facility-override <i>facility-name</i>;
 interface-service <i>prefix-value</i>;
 }
}</pre> |
| Hierarchy Level | [edit services service-set <i>service-set-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure generation of system log messages for the service set. The system log information is passed to the kernel for logging in the <code>/var/log</code> directory. These settings override the values defined at the [edit interfaces <i>interface-name</i> services-options] hierarchy level; for more information on configuring those values, see “ Configuring System Logging for Services Interfaces ” on page 20. |
| Options | The remaining statements are described separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring System Logging for Service Sets on page 47 |

syslog (Services Stateful Firewall)

| | |
|---------------------------------|--|
| Syntax | syslog; |
| Hierarchy Level | [edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Enable system logging. The system log information from the Adaptive Services or Multiservices PIC is passed to the kernel for logging in the /var/log directory. This setting overrides any syslog statement setting included in the service set or interface default configuration. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Actions in Stateful Firewall Rules on page 334 |

syn-cookie

| | |
|---------------------------------|--|
| Syntax | <pre>syn-cookie {
 mss value;
 threshold rate;
}</pre> |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>Enable SYN-cookie defenses against SYN attacks. By default, SYN-cookie techniques are not applied.</p> <p>When SYN cookie is enabled on Junos OS and becomes the TCP-negotiating proxy for the destination server, it replies to each incoming SYN segment with a SYN/ACK containing an encrypted cookie as its initial sequence number (ISN). The cookie is an MD5 hash of the original source address and port number, destination address and port number, and ISN from the original SYN packet. After sending the cookie, Junos OS drops the original SYN packet and deletes the calculated cookie from memory. If there is no response to the packet containing the cookie, the attack is noted as an active SYN attack and is effectively stopped.</p> <p>If the initiating host responds with a TCP packet containing the cookie +1 in the TCP ACK field, Junos OS extracts the cookie, subtracts 1 from the value, and recomputes the cookie to validate that it is a legitimate ACK. If it is legitimate, Junos OS starts the TCP proxy process by setting up a session and sending a SYN to the server containing the source information from the original SYN. When Junos OS receives a SYN/ACK from the server, it sends ACKs to the server and to the initiation host. At this point the connection is established and the host and server are able to communicate directly.</p> |
| Options | The remaining statements are described separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Actions in IDS Rules on page 358 |

tcp-mss

| | |
|---------------------------------|---|
| Syntax | <code>tcp-mss <i>number</i>;</code> |
| Hierarchy Level | [edit services service-set service-set-name] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Specify the TCP Maximum Segment Size (MSS) allowed for the service set. |
| Options | <i>number</i> —MSS value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Set Limitations on page 37 |

term (Services CoS)

Syntax `term term-name {`
 `from {`
 `application-sets set-name;`
 `applications [application-names];`
 `destination-address address;`
 `destination-prefix-list list-name <except>;`
 `source-address address;`
 `source-prefix-list list-name <except>;`
 `}`
 `then {`
 `application-profile profile-name;`
 `dscp (alias | bits);`
 `forwarding-class class-name;`
 `syslog;`
 `(reflexive | reverse) {`
 `application-profile profile-name;`
 `dscp (alias | bits);`
 `forwarding-class class-name;`
 `syslog;`
 `}`
 `}`
 `}`

Hierarchy Level [edit [services](#) cos [rule](#) *rule-name*]

Release Information Statement introduced in Junos OS Release 8.1.

Description Define the CoS term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Configuring CoS Rules on page 514](#).

term (Services IDS)

```
Syntax  term term-name {
        from {
            application-sets set-name;
            applications [ application-names ];
            destination-address (address | any-unicast) <except>;
            destination-address-range low minimum-value high maximum-value <except>;
            source-address (address | any-unicast) <except>;
            source-address-range low minimum-value high maximum-value <except>;
        }
        then {
            aggregation {
                destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
                source-prefix prefix-value | source-prefix-ipv6 prefix-value;
            }
            (force-entry | ignore-entry);
            logging {
                syslog;
                threshold rate;
            }
            session-limit {
                by-destination {
                    hold-time seconds;
                    maximum number;
                    packets number;
                    rate number;
                }
                by-pair {
                    hold-time seconds;
                    maximum number;
                    packets number;
                    rate number;
                }
                by-source {
                    hold-time seconds;
                    maximum number;
                    packets number;
                    rate number;
                }
            }
            syn-cookie {
                mss value;
                threshold rate;
            }
        }
    }
```

Hierarchy Level [edit [services](#) ids [rule](#) *rule-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the IDS term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related • [Configuring IDS Rules on page 355](#)
Documentation

term (Services IPsec VPN)

```
Syntax  term term-name {
        from {
            destination-address address;
            ipsec-inside-interface interface-name;
            source-address address;
        }
        then {
            anti-replay-window-size bits;
            backup-remote-gateway address;
            clear-dont-fragment-bit;
            dynamic {
                ike-policy policy-name;
                ipsec-policy policy-name;
            }
            initiate-dead-peer-detection;
            manual {
                direction (inbound | outbound | bidirectional) {
                    authentication {
                        algorithm (hmac-md5-96 | hmac-sha1-96);
                        key (ascii-text key | hexadecimal key);
                    }
                    auxiliary-spi spi-value;
                    encryption {
                        algorithm algorithm;
                        key (ascii-text key | hexadecimal key);
                    }
                    protocol (ah | bundle | esp);
                    spi spi-value;
                }
            }
            no-anti-replay;
            remote-gateway address;
            syslog;
            tunnel-mtu bytes;
        }
    }
```

Hierarchy Level [edit [services](#) ipsec-vpn [rule](#) *rule-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the IPsec term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Required Privilege interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

Related Documentation • [Configuring IPsec Rules on page 422](#)

term (Services HCM)

Syntax

```
term term-num {  
  from {  
    url-list url-list-name;  
    url url_identifier {  
      host hostname;  
      request-url page-name;  
    }  
  }  
}
```

Hierarchy Level [edit services hcm url-rule *url-rule-name* term *term-num*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify a numbered identity for each term inside a rule.

Options *term-num*—Identifier value for the term.

Range: 1 through 255

Default: If no value is entered, the default value is 1.

Required Privilege interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

term (Services NAT)

```
Syntax  term term-name {
        from {
            application-sets set-name;
            applications [ application-names ];
            destination-address (address | any-unicast) <except>;
            destination-address-range low minimum-value high maximum-value <except>;
            source-address (address | any-unicast) <except>;
            source-address-range low minimum-value high maximum-value <except>;
        }
        then {
            no-translation;
            translated {
                address-pooling paired;
                destination-pool nat-pool-name;
                destination-prefix destination-prefix;
                dns-alg-pool dns-alg-pool;
                dns-alg-prefix dns-alg-prefix;
                filtering-type endpoint-independent;
                mapping-type endpoint-independent;
                source-pool nat-pool-name;
                source-prefix source-prefix;
                translation-type (basic-nat-pt | basic-nat44 | basic-nat66 | dnat-44 | dynamic-nat44
                    | napt-44 | napt-66 | napt-pt | stateful-nat64 | twice-basic-nat-44 |
                    twice-dynamic-nat-44 | twice-napt-44);
            }
        }
        syslog;
    }
```

Hierarchy Level [edit [services](#) nat [rule](#) *rule-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the NAT term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Network Address Translation Rules Overview on page 69](#)

term (Services Stateful Firewall)

Syntax `term term-name {
 from {
 application-sets set-name;
 applications [application-names];
 destination-address (address | any-unicast) <except>;
 destination-address-range low minimum-value high maximum-value <except>;
 destination-prefix-list list-name <except>;
 source-address (address | any-unicast) <except>;
 source-address-range low minimum-value high maximum-value <except>;
 source-prefix-list list-name <except>;
 }
 then {
 (accept | discard | reject);
 syslog;
 }
 }`

Hierarchy Level [edit [services](#) stateful-firewall [rule](#) *rule-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the stateful firewall term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Usage Guidelines See “[Configuring Stateful Firewall Rules](#)” on page 331.

Required Privilege interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

then (Services CoS)

| | |
|---------------------------------|---|
| Syntax | <pre> then { application-profile profile-name; dscp (alias bits); forwarding-class class-name; syslog; (reflexive reverse) { application-profile profile-name; dscp (alias bits); forwarding-class class-name; syslog; } } </pre> |
| Hierarchy Level | [edit services cos rule rule-name term term-name] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | <p>Define the CoS term actions.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> Configuring Actions in a CoS Rule Configuring Actions in CoS Rules on page 516 |

then (Services HCM)

| | |
|---------------------------------|--|
| Syntax | <pre> then { discard; accept; count; log-request; } </pre> |
| Hierarchy Level | [edit services hcm url-rule url-rule-name term term-num] |
| Release Information | Statement introduced in Junos OS Release 12.1. |
| Description | Define the HCM term actions. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |

then (Services IDS)

```
Syntax  then {
        aggregation {
            destination-prefix prefix-number | destination-prefix-ipv6 prefix-value;
            source-prefix prefix-number | source-prefix-ipv6 prefix-value;
        }
        (force-entry | ignore-entry);
        logging {
            syslog;
            threshold rate;
        }
        session-limit {
            by-destination {
                hold-time seconds;
                maximum number;
                packets number;
                rate number;
            }
            by-pair {
                hold-time seconds;
                maximum number;
                packets number;
                rate number;
            }
            by-source {
                hold-time seconds;
                maximum number;
                packets number;
                rate number;
            }
        }
        syn-cookie {
            mss value;
            threshold rate;
        }
    }
```

Hierarchy Level [edit [services](#) ids [rule](#) *rule-name* [term](#) *term-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the IDS term actions.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring IDS Rules on page 355](#)

then (Services IPsec VPN)

```
Syntax  then {
        anti-replay-window-size bits;
        backup-remote-gateway address;
        clear-dont-fragment-bit;
        dynamic {
            ike-policy policy-name;
            ipsec-policy policy-name;
        }
        initiate-dead-peer-detection;
        dead-peer-detection {
            interval seconds;
            threshold number;
        }
        manual {
            direction (inbound | outbound | bidirectional) {
                authentication {
                    algorithm (hmac-md5-96 | hmac-sha1-96);
                    key (ascii-text key | hexadecimal key);
                }
                auxiliary-spi spi-value;
                encryption {
                    algorithm algorithm;
                    key (ascii-text key | hexadecimal key);
                }
                protocol (ah | bundle | esp);
                spi spi-value;
            }
        }
        no-anti-replay;
        remote-gateway address;
        syslog;
        tunnel-mtu bytes;
    }
```

Hierarchy Level [edit [services](#) ipsec-vpn [rule](#) *rule-name* [term](#) *term-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the IPsec term actions.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring IPsec Rules on page 422](#)

then (Services NAT)

Syntax then {
 no-translation;
 translated {
 address-pooling paired;
 destination-pool *nat-pool-name*;
 destination-prefix (Services NAT) *destination-prefix*;
 dns-alg-pool *dns-alg-pool*;
 dns-alg-prefix *dns-alg-prefix*;
 filtering-type endpoint-independent;
 mapping-type endpoint-independent;
 source-pool *nat-pool-name*;
 source-prefix *source-prefix*;
 translation-type (basic-nat-pt | basic-nat44 | basic-nat66 | dnat-44 | dynamic-nat44
 | napt-44 | napt-66 | napt-pt | stateful-nat64 | twice-basic-nat-44 |
 twice-dynamic-nat-44 | twice-napt-44);
 }
 }
 syslog;
 }

Hierarchy Level [edit [services](#) nat [rule](#) *rule-name* [term](#) *term-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the NAT term actions.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Network Address Translation Rules Overview on page 69](#)

then (Services Stateful Firewall)

| | |
|---------------------------------|---|
| Syntax | <pre>then { (accept discard reject); syslog; }</pre> |
| Hierarchy Level | [edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the stateful firewall term actions. You can configure the router to accept, discard, or reject the targeted traffic. The other actions are optional. |
| Options | <p>accept—Accept the traffic and send it on to its destination.</p> <p>discard—Do not accept traffic or process it further.</p> <p>reject—Do not accept the traffic and return a rejection message. Rejected traffic can be logged or sampled.</p> <p>The remaining statement is explained separately.</p> |
| Usage Guidelines | See “Configuring Actions in Stateful Firewall Rules” on page 334 . |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> |

threshold (Services IPsec)

| | |
|---------------------------------|---|
| Syntax | <code>threshold <i>number</i>;</code> |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then dead-peer-detection] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Specify the maximum number of unsuccessful dead peer detection (DPD) requests to be sent before the peer is considered unavailable. (The threshold value is used for IKEv1 security associations (SAs) but not for IKEv2 SAs.) |
| Options | number —Maximum number of unsuccessful DPD requests to be sent.
Range: 1 through 10
Default: 3 |
| Required Privilege Level | security—To view this statement in the configuration.
security-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IPsec Rules on page 422 |

threshold (Services Logging and SYN-Cookie Defenses)

| | |
|---------------------------------|--|
| Syntax | <code>threshold <i>rate</i>;</code> |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> then logging],
[edit services ids rule <i>rule-name</i> term <i>term-name</i> then syn-cookie] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the threshold for logging or applying SYN-cookie defenses. |
| Options | rate —Logging threshold number of events per second.
rate —SYN-cookie defense number of SYN attacks per second. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Actions in IDS Rules on page 358 |

traceoptions (Security PKI)

| | |
|------------------------|--|
| Syntax | <pre> traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i>> <size <i>maximum-file-size</i>> <world-readable no-world-readable>; flag <i>flag</i>; } </pre> |
| Hierarchy Level | [edit security pki] |
| Description | Configure security public key infrastructure (PKI) trace options. To specify more than one trace option, include multiple flag statements. Trace option output is recorded in the <code>/var/log/pkid</code> file. |
| Options | <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. To include the file statement, you must specify a filename.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file (for example, pkid) reaches its maximum size, it is renamed pkid.0, then pkid.1, and so on, until the maximum number of trace files is reached. When the maximum number is reached, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>flag—Trace operation to perform. To specify more than one trace operation, include multiple flag statements:</p> <ul style="list-style-type: none"> all—Trace with all flags enabled. certificate-verification—Trace PKI certificate verification events. online-crl-check—Trace PKI online certificate revocation list (CRL) events. enrollment—PKI certificate enrollment tracing. <p>match <i>regular-expression</i>—(Optional) Refine the output to include lines that contain the regular expression.</p> <p>size <i>maximum-file-size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB). If you specify a maximum file size, you also must specify a maximum number of trace files with the files <i>number</i> option.</p> <p>Default: 1024 KB</p> <p>world-readable no-world-readable—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The world-readable option enables any user to read the file. To explicitly set the default behavior, use the no-world-readable option.</p> |

| | |
|------------------------------|--|
| Required Privilege | trace—To view this statement in the configuration. |
| Level | trace-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Tracing Junos VPN Site Secure Operations on page 436 |

traceoptions (Services IPsec VPN)

| | |
|----------------------------|---|
| Syntax | <pre> traceoptions { file <filename> <files number> <match regular-expression> <size bytes> <world-readable no-world-readable>; flag flag; level level; no-remote-trace; } </pre> |
| Hierarchy Level | [edit services ipsec-vpn] |
| Release Information | Statement introduced in Junos OS Release 7.5.
level option added in Junos OS Release 10.0. |
| Description | Configure IPsec tracing operations. By default, messages are written to <code>/var/log/kmd</code> . |
| Options | <p>files <i>number</i>—Maximum number of trace data files.
 Range: 2 through 1000</p> <p>flag <i>flag</i>—Tracing operation to perform:</p> <ul style="list-style-type: none"> • all—Trace everything. • certificates—Trace certificates that apply to the IPsec service set. • database—Trace security associations database events. • general—Trace general events. • ike—Trace IKE module processing. • parse—Trace configuration processing. • policy-manager—Trace policy manager processing. • routing-socket—Trace routing socket messages. • snmp—Trace SNMP operations. • timer—Trace internal timer events. <p>level <i>level</i>—Key management process (kmd) tracing level. The following values are supported:</p> <ul style="list-style-type: none"> • all—Match all levels. • error—Match error conditions. • info—Match informational messages. • notice—Match conditions that should be handled specially. • verbose—Match verbose messages. • warning—Match warning messages. |

size bytes—Maximum trace file size.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation • [Tracing Junos VPN Site Secure Operations on page 436](#)

traceoptions (Services L2TP)

| | |
|----------------------------|--|
| Syntax | <pre> traceoptions { debug-level <i>level</i>; file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i> > <size <i>maximum-file-size</i>> <world-readable no-world-readable>; filter { protocol <i>name</i>; user <i>user@domain</i>; user-name <i>username</i>; } flag <i>flag</i>; interfaces <i>interface-name</i> { debug-level <i>level</i>; flag <i>flag</i>; } level (all error info notice verbose warning); no-remote-trace; } </pre> |
| Hierarchy Level | [edit services l2tp] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define tracing operations for L2TP processes. |
| Options | <p>debug-level <i>level</i>—Trace level for PPP, L2TP, RADIUS, and UDP; this option does not apply to L2TP on MX Series routers:</p> <ul style="list-style-type: none"> detail—Trace detailed debug information. error—Trace error information. packet-dump—Trace packet decoding information. <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 3 files</p> <p>filter—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifiers simplifies troubleshooting in a scaled environment.</p> <ul style="list-style-type: none"> protocol <i>name</i>—One of the following protocols; this option does not apply to L2TP on MX Series routers: <ul style="list-style-type: none"> l2tp |

- **ppp**
- **radius**
- **udp**
- **user** *user@domain*—Username of a subscriber; this option does not apply to L2TP on M Series routers. Optionally use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms.
- **user-name** *username*—Username of a subscriber; this option does not apply to L2TP on MX Series routers.

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.
- **configuration**—Trace configuration events.
- **events**—Trace interface events.
- **general**—Trace general events.
- **gres**—Trace GRES events.
- **init**—Trace daemon initialization.
- **ipc-rx**—Trace IPC receive events.
- **ipc-tx**—Trace IPC transmit events.
- **memory**—Trace memory management code.
- **message**—Trace message processing code.
- **packet-error**—Trace packet error events.
- **parse**—Trace parsing events.
- **protocol**—Trace L2TP events.
- **receive-packets**—Trace received L2TP packets.
- **routing-process**—Trace routing process interactions.
- **routing-socket**—Trace routing socket events.
- **session-db**—Trace session database interactions.
- **states**—Trace state machine events.
- **timer**—Trace timer events.
- **transmit-packets**—Trace transmitted L2TP packets.
- **tunnel**—Trace tunnel events.

interfaces *interface-name*—Apply L2TP traceoptions to a specific services interface. This option does not apply to L2TP on MX Series routers.

- **debug-level *level***—Trace level for the interface; this option does not apply to L2TP on MX Series routers:
 - **detail**—Trace detailed debug information.
 - **error**—Trace error information.
 - **extensive**—Trace all PIC debug information.
- **flag *flag***—Tracing operation to perform for the interface. This option does not apply to L2TP on MX Series routers. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:
 - **all**—Trace everything.
 - **ipc**—Trace L2TP Inter-Process Communication (IPC) messages between the PIC and the Routing Engine.
 - **packet-dump**—Dump each packet content based on debug level.
 - **protocol**—Trace L2TP, PPP, and multilink handling.
 - **system**—Trace packet processing on the PIC.

level—Specify level of tracing to perform. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- **all**—Match messages of all levels.
- **error**—Match error messages.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure **verbose**, messages at all higher levels are traced. Therefore, the result is the same as when you configure **all**.
- **warning**—Match warning messages.

Default: error

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

size *maximum-file-size*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB

Range: 10240 through 1073741824

world-readable—(Optional) Enable unrestricted file access.

| | |
|---------------------------------|--|
| Required Privilege Level | trace—To view this statement in the configuration. |
| | trace-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Tracing L2TP Operations on page 652• <i>Tracing L2TP Operations for Subscriber Access</i> |

traceoptions (Services Logging)

| | |
|----------------------------|---|
| Syntax | <pre> traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; no-remote-trace; } </pre> |
| Hierarchy Level | [edit services adaptive-services-pics],
[edit services logging] |
| Release Information | Statement introduced before Junos OS Release 7.4.
file option added in Release 8.0. |
| Description | Configure Adaptive Services or Multiservices PIC tracing operations. The messages are output to /var/log/serviced . |
| Options | <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. All files are placed in the directory /var/log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 3 files</p> <p>flag <i>flag</i>—Tracing operation to perform:</p> <ul style="list-style-type: none"> • all—Trace everything. • command-queued—Trace command enqueue events. • config—Trace configuration events. • handshake—Trace handshake events. • init—Trace initialization events. • interfaces—Trace interface events. • mib—Trace GGSN SNMP MIB events. • removed-client—Trace client cleanup events. • show—Trace CLI command servicing. <p>match <i>regex</i>—(Optional) Match output to a defined regular expression (regex).</p> |

Default: If you do not include this option, the trace operation output includes all lines relevant to the logged events.

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

| | |
|---------------------------------|--|
| Required Privilege Level | interface—To view this statement in the configuration. |
| | interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Tracing Services PIC Operations on page 48 |

translated

| | |
|---------------------------------|---|
| Syntax | <pre>translated { address-pooling paired; destination-pool nat-pool-name; destination-prefix destination-prefix; dns-alg-pool dns-alg-pool; dns-alg-prefix dns-alg-prefix; filtering-type endpoint-independent; mapping-type endpoint-independent; overload-pool overload-pool-name; overload-prefix; source-pool nat-pool-name; translation-type (basic-nat-pt basic-nat44 basic-nat66 dnat-44 dynamic-nat44 napt-44 napt-66 napt-pt stateful-nat64 twice-basic-nat-44 twice-dynamic-nat-44 twice-napt-44) }</pre> |
| Hierarchy Level | [edit services nat rule rule-name term term-name then] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define properties for translated traffic. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Network Address Translation Rules Overview on page 69 |

trigger-link-failure

| | |
|---------------------------------|--|
| Syntax | trigger-link-failure <i>interface-name</i> ; |
| Hierarchy Level | [edit interfaces <i>lsq-fpc/pic/port</i> lsq-failure-options] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | List of SONET interfaces connected to the LSQ interface that can implement Automatic Protection Switching (APS) if the Link Services IQ PIC fails. |
| Options | <i>interface-name</i> —Name of SONET interface. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Association between LSQ and SONET Interfaces on page 546 |

translated-port

| | |
|---------------------------------|---|
| Syntax | <code>translated-port <i>port id</i>;</code> |
| Hierarchy Level | [edit services nat port-forwarding <i>map-name</i>] |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Specify the port to which all traffic will be translated. |
| Options | <i>port id</i> —The port number to which traffic will be translated. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• port-forwarding on page 1433• destined-port on page 1345 |

translation-type

| | |
|----------------------------|--|
| Syntax | translation-type (basic-nat-pt basic-nat44 basic-nat66 nat-44 deterministic-napt44 dnat-44 dynamic-nat44 napt-44 napt-66 napt-pt nptv6 stateful-nat64 twice-basic-nat-44 twice-dynamic-nat-44 twice-napt-44) |
| Hierarchy Level | [edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated] |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>The following options introduced in Junos OS Release 11.2, replacing all previous options:</p> <ul style="list-style-type: none"> • basic-nat44 • basic-nat66 • basic-nat-pt • deterministic-napt44 • dnat-44 • dynamic-nat44 • napt-44 • napt-66 • napt-pt • stateful-nat64 <p>twice-basic-nat-44 option introduced in Junos OS Release 11.4.</p> <p>twice-dynamic-nat-44 option introduced in Junos OS Release 11.4.</p> <p>twice-napt-44 option introduced in Junos OS Release 11.4.</p> <p>deterministic-napt44 option introduced in Junos OS Release 12.1.</p> <p>nptv6 option introduced in Junos OS Release 15.1</p> |
| Description | Specify the NAT translation types. |
| Options | <ul style="list-style-type: none"> • basic-nat44—Translate the source address statically (IPv4 to IPv4). • basic-nat66—Translate the source address statically (IPv6 to IPv6). • basic-nat-pt—Translate the addresses of IPv6 hosts as they originate sessions to the IPv4 hosts in the external domain. The basic-nat-pt option is always implemented with DNS ALG. • deterministic-napt44—Translate as napt-44, and use deterministic port block allocation for port translation. • dnat-44—Translate the destination address statically (IPv4 to IPv4). • dynamic-nat44—Translate only the source address by dynamically choosing the NAT address from the source address pool. |

- **napt-44**—Translate the transport identifier of the IPv4 private network to a single IPv4 external address.
- **napt-66**—Translate the transport identifier of the IPv6 private network to a single IPv6 external address.
- **napt-pt**—Bind addresses in an IPv6 network with addresses in an IPv4 network and vice versa to provide transparent routing for the datagrams traversing between the address realms.
- **nptv6**—Translate the source address prefix in a stateless manner (IPv6 to IPv6).
- **stateful-nat64**—Implement dynamic address and port translation for source IP addresses (IPv6-to-IPv4) and prefix removal translation for the destination IP addresses (IPv6-to-IPv4).
- **twice-basic-nat-44**—Translate the source and destination addresses statically (IPv4 to IPv4).
- **twice-dynamic-nat-44**—Translate the source address by dynamically choosing the NAT address from the source address pool. Translate the destination address statically.
- **twice-dynamic-napt-44**—Translate the transport identifier of the IPv4 private network to a single IPv4 external address. Translate the destination address statically.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Network Address Translation Rules Overview on page 69](#)

trusted-ca

Syntax `trusted-ca ca-profile-name;`

Hierarchy Level [edit [services service-set](#) *service-set-name* [ipsec-vpn-options](#)]

Release Information Statement introduced in Junos OS Release 7.5.

Description Identify one or more trusted IPsec certification authorities.

Options *ca-profile-name*—Name of certification authority profile, which is configured at the [edit [security pki](#)] hierarchy level.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring IPsec Service Sets on page 430](#)

ttl-threshold

| | |
|---------------------------------|--|
| Syntax | <code>ttl-threshold <i>number</i>;</code> |
| Hierarchy Level | [edit applications application application-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the traceroute time-to-live (TTL) threshold value. This value sets the acceptable level of network penetration for trace routing. |
| Options | <i>number</i> —TTL threshold value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• ALG Descriptions on page 277• Configuring the TTL Threshold on page 320.• Examples: Configuring Application Protocols on page 321• Verifying the Output of ALG Sessions |

tunnel-group

Syntax `tunnel-group group-name {`
 `aaa-access-profile profile-name;`
 `dynamic-profile profile-name;`
 `hello-interval seconds;`
 `hide-avps;`
 `l2tp-access-profile profile-name;`
 `local-gateway address {`
 `address address;`
 `gateway-name gateway-name;`
 `}`
 `maximum-send-window packets;`
 `ppp-access-profile profile-name;`
 `receive-window packets;`
 `retransmit-interval seconds;`
 `service-device-pool pool-name;`
 `service-interface interface-name;`
 `syslog {`
 `host hostname {`
 `services severity-level;`
 `facility-override facility-name;`
 `log-prefix prefix-value;`
 `}`
 `}`
 `tos-reflect;`
 `tunnel-switch-profile profile-name;`
 `tunnel-timeout seconds;`
 `}`

Hierarchy Level [edit services l2tp]

Release Information Statement introduced before Junos OS Release 7.4.
 Support for MX Series routers introduced in Junos OS Release 11.4.

Description Specify the L2TP tunnel properties. On MX Series routers, you can configure up to 256 tunnel groups. On M Series routers, there is no limit to the number of tunnel groups you can configure.



NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options *group-name*—Identifier for the tunnel group.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring L2TP Tunnel Groups on page 641](#)
 - [Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces](#)

tunnel-mtu (Services IPsec VPN)


| | |
|----------------------------|---|
| Syntax | tunnel-mtu <i>bytes</i> ; |
| Hierarchy Level | [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced in Junos OS Release 7.5. |
| Description | Maximum transmission unit (MTU) size for IPsec tunnels. |
| Options | <p><i>bytes</i>—MTU size.</p> <p>Default: 1500 bytes</p> <p>Range: 256 through 9192 bytes</p> |



NOTE: Clear the IPsec SA in tunnel-mtu to accommodate Jumbo frames larger than 1500 bytes.

| | |
|---------------------------------|--|
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Specifying the MTU for IPsec Tunnels • mtu on page 1611 |

tunnel-mtu (Services Service Set)

| | |
|---------------------------------|---|
| Syntax | <code>tunnel-mtu bytes;</code> |
| Hierarchy Level | [edit services service-set <i>service-set-name</i> ipsec-vpn-options] |
| Release Information | Statement introduced in Junos OS Release 10.0. |
| Description | <p>Maximum transmission unit (MTU) size for IPsec tunnels. This statement is useful for dynamic endpoint tunnels for which you cannot configure the tunnel-mtu statement at the [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] hierarchy level.</p> <p>For static IPsec tunnels, this statement sets the tunnel MTU value for all the tunnels within this service set. If you need a specific value for a particular tunnel, then set the tunnel-mtu statement at the [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] hierarchy level.</p> |
| | <div> NOTE: The tunnel-mtu setting at the [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then] hierarchy level overrides the value specified at the [edit services service-set <i>service-set-name</i> ipsec-vpn-options] hierarchy level.</div> |
| Options | <p><i>bytes</i>—MTU size.</p> <p>Default: 1500 bytes</p> <p>Range: 256 through 9192 bytes</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• mtu on page 1611• Configuring IPsec Service Sets on page 430• Specifying the MTU for IPsec Tunnels |

tunnel-timeout

| | |
|---------------------------------|--|
| Syntax | tunnel-timeout <i>seconds</i> ; |
| Hierarchy Level | [edit services l2tp tunnel-group <i>name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the maximum downtime for an L2TP tunnel, after which the tunnel is terminated because the connection is presumed to have been lost. |
| Options | seconds —Interval after which the tunnel is terminated if no data can be sent.
Default: 120 seconds |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Timers for L2TP Tunnels on page 643 • Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces |

url

| | |
|---------------------------------|--|
| Syntax | url <i>url_identifier</i> ; |
| Hierarchy Level | [edit services hcm url-rule <i>url-rule-name</i> term <i>term-num</i> from] |
| Release Information | Statement introduced in Junos OS Release 12.1. |
| Description | Specify an integer that uniquely identifies a particular URL definition within a term. |
| Options | url_identifier —URL identifier number.
Range: 1 through 32,767
Default: If no value is added, the default value is 1. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

url-list

| | |
|---------------------------------|--|
| Syntax | <code>url-list <i>url-list-name</i> ;</code> |
| Hierarchy Level | <code>[edit services hcm url-rule <i>url-rule-name</i> term <i>term-num</i> from]</code> |
| Release Information | Statement introduced in Junos OS Release 12.1. |
| Description | Specify the name of a previously defined URL list to be included as a matching condition. A match for the term is considered when a URL matches any hostname and any request-URL within the same term. |
| Options | <i>url-list-name</i> —Name of the previously defined URL list. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

url-rule

| | |
|---------------------------------|---|
| Syntax | <pre>url-rule <i>url-rule-name</i> {
 term <i>term-num</i> {
 from {
 url-list <i>url-list-name</i>;
 url <i>url_identifier</i> {
 host <i>hostname</i>;
 request-url <i>page-name</i>;
 }
 }
 then {
 discard;
 accept;
 count;
 log-request;
 }
 }
}</pre> |
| Hierarchy Level | <code>[edit services hcm]</code> |
| Release Information | Statement introduced in Junos OS Release 12.1. |
| Description | Specify the name of the URL rule. |
| Options | <i>url-rule-name</i> —Name of the URL rule. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

url-rule-set

| | |
|---------------------------------|--|
| Syntax | <code>url-rule-set <i>url-rule-set-name</i> {
 url-rule <i>rule1</i>;
 url-rule <i>rule2</i>;
}</code> |
| Hierarchy Level | <code>[edit services hcm url-rule <i>url-rule-name</i>]</code> |
| Release Information | Statement introduced in Junos OS Release 11.2. |
| Description | Specify the name of the rule set. A rule set is a collection of rules ordered in the sequence in which they are entered. |
| Options | <i>url-rule-set-name</i> —Name of the collection of URL rules that constitute this rule set. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

unit (Aggregated Multiservices)

| | |
|----------------------------|---|
| Syntax | <code>unit <i>interface-unit-number</i> {
 family <i>family</i>;
}</code> |
| Hierarchy Level | <code>[edit interfaces <i>interface-name</i>]</code> |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Configure the logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

The remaining statements are explained separately. |
| Options | <i>interface-unit-number</i> —Number of the logical unit. |



NOTE: Unit 0 is reserved and cannot be configured under the aggregated Multiservices interface (ams).

Range: 1 through 16,384

| | |
|---------------------------------|--|
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Understanding Aggregated Multiservices Interfaces on page 599 • Example: Configuring an Aggregated Multiservices Interface (AMS) on page 608 • interfaces on page 1386 |

unit (Interfaces)

| | |
|---------------------------------|---|
| Syntax | <pre>unit <i>logical-unit-number</i> {
 family inet {
 address <i>address</i> {
 }
 service {
 input {
 [<i>service-set</i> <i>service-set-name</i> <<i>service-filter</i> <i>filter-name</i>>];
 <i>post-service-filter</i> <i>filter-name</i>;
 }
 output {
 [<i>service-set</i> <i>service-set-name</i> <<i>service-filter</i> <i>filter-name</i>>];
 }
 }
 <i>service-domain</i> (inside outside);
 }
}</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device. |
| Options | <p><i>logical-unit-number</i>—Number of the logical unit.</p> <p>Range: 0 through 16,384</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces. |

unit (Voice Services)

```
Syntax  unit logical-unit-number {
        compression {
            rtp {
                f-max-period number;
                maximum-contexts number <force>;
            }
            port {
                minimum port-number;
                maximum port-number;
            }
            queues [ queue-numbers ];
        }
        compression-device interface-name;
        encapsulation type;
        family family {
            address address {
                ...
            }
            bundle (lsq-fpc/pic/port | ...);
        }
    }
```

Hierarchy Level [edit [interfaces](#) *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options *logical-unit-number*—Number of the logical unit.

Range: 0 through 16,384

The remaining statements are explained separately.

Required Privilege interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Network Interfaces Library for Routing Devices* for other statements that do not affect services interfaces.

- [Configuring Services Interfaces for Voice Services on page 622](#)


uuid

| | |
|---------------------------------|---|
| Syntax | <code>uuid <i>hex-value</i>;</code> |
| Hierarchy Level | [edit applications application application-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the Universal Unique Identifier (UUID) for DCE RPC objects. |
| Options | <i>hex-value</i> —Hexadecimal value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• ALG Descriptions on page 277• Configuring a Universal Unique Identifier on page 320• Examples: Configuring Application Protocols on page 321• Verifying the Output of ALG Sessions |

v6rd

| | |
|---------------------------------|--|
| Syntax | <pre>v6rd v6rd-softwire-concentrator { ipv4-prefix <i>ipv4-prefix</i>; v6rd-prefix <i>ipv6-prefix</i>; mtu-v4 <i>mtu-v4</i>; softwire-address <i>ipv4-address</i>; }</pre> |
| Hierarchy Level | [edit services softwire softwire-concentrator] |
| Release Information | Statement introduced in Junos OS Release 10.4. |
| Description | Configure settings for a 6rd concentrator used to process IPv6 packets encapsulated in IPv4 packets. |
| Options | <p><i>ipv4-prefix</i>—IPv4 prefix of the customer edge (CE) network</p> <p><i>ipv6-prefix</i>—IPv6 prefix of the 6rd domain.</p> <p><i>mtu-v4</i>—Maximum transmission unit (MTU), in bytes (576 through 9192), for IPv6 packets encapsulated into IPv4. If the final length is greater than the configured value, the IPv4 packet will be dropped.</p> <p><i>address</i>—IPv4 address of a softwire concentrator. This is an IPv4 address independent of any interface and on a different prefix.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring a 6rd Softwire Concentrator on page 245 |

version (IKE)

| | |
|---|---|
| Syntax | version (1 2); |
| Hierarchy Level | [edit services ipsec-vpn ike policy <i>policy-name</i>], |
| Release Information | Statement introduced in Junos OS Release 11.4. |
| Description | Configure the Internet Key Exchange (IKE) version that is used to negotiate dynamic SAs for IPSec. |
| Options | 1—Uses IKEv1.
2—Uses IKEv2. |
| <hr/> | |
| <div> NOTE: By default, Junos OS uses IKE policy version 1.0. Version 2.0 is supported only in Junos OS Release 11.4 and later. If no version is explicitly configured, Junos OS sets the version to version 1.0.</div> <hr/> | |
| Required Privilege Level | admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring IKE Policies on page 409 |

video (Application Profile)

| | |
|---------------------------------|---|
| Syntax | video {
dscp (<i>alias</i> <i>bits</i>);
forwarding-class <i>class-name</i> ;
} |
| Hierarchy Level | [edit services cos application-profile <i>profile-name</i> sip] |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | Set the appropriate dscp and forwarding-class values for SIP video traffic. |
| Default | By default, the system will not alter the DSCP or forwarding class for SIP video traffic. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Application Profiles• voice (Application Profile) on page 1533 |

voice (Application Profile)

| | |
|---------------------------------|--|
| Syntax | voice {
dscp (<i>alias</i> <i>bits</i>);
forwarding-class <i>class-name</i> ;
} |
| Hierarchy Level | [edit services cos application-profile <i>profile-name</i> sip] |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | Set the appropriate dscp and forwarding-class values for SIP voice traffic. |
| Default | By default, the system will not alter the DSCP or forwarding class for SIP voice traffic. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Application Profiles • video (Application Profile) on page 1532 |

warm-standby

| | |
|---------------------------------|---|
| Syntax | warm-standby; |
| Hierarchy Level | [edit interfaces <i>rls</i> <i>number</i> redundancy-options] |
| Release Information | Statement introduced in Junos OS Release 8.0. |
| Description | For AS or Multiservices PIC redundancy configurations, specify that the failure detection and recovery involves one backup PIC supporting multiple working PICs. Recovery time is not guaranteed. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 548 |

Application Aware Services Configuration Statements

- [\[edit services aacl\] Hierarchy List on page 1536](#)
- [\[edit services application-identification\] Hierarchy Level on page 1537](#)
- [aacl-fields on page 1539](#)
- [aacl-statistics-profile on page 1540](#)
- [address on page 1541](#)

- [application \(Defining\) on page 1542](#)
- [application \(Including in Rule\) on page 1543](#)
- [application-aware-access-list-fields on page 1544](#)
- [application-group on page 1545](#)
- [application-group-any on page 1545](#)
- [application-groups \(Services AAACL\) on page 1546](#)
- [application-groups \(Services Application Identification\) on page 1546](#)
- [application-system-cache-timeout on page 1547](#)
- [application-unknown on page 1547](#)
- [applications \(Services AAACL\) on page 1547](#)
- [applications \(Services Application Identification\) on page 1548](#)
- [automatic on page 1548](#)
- [bypass-traffic-on-exceeding-flow-limits on page 1549](#)
- [chain-order on page 1549](#)
- [context on page 1550](#)
- [destination \(Services\) on page 1550](#)
- [destination-address on page 1551](#)
- [destination-address-range on page 1551](#)
- [destination-prefix-list on page 1552](#)
- [direction on page 1552](#)
- [disable \(APPID Application\) on page 1553](#)
- [disable \(APPID Application Group\) on page 1553](#)
- [disable \(APPID Port Mapping\) on page 1553](#)
- [disable-global-timeout-override on page 1554](#)
- [download on page 1554](#)
- [enable-asymmetric-traffic-processing on page 1555](#)
- [enable-heuristics on page 1555](#)
- [file on page 1556](#)
- [from on page 1557](#)
- [idle-timeout on page 1558](#)
- [ignore-errors on page 1558](#)
- [index \(Applications\) on page 1559](#)
- [index \(Nested Applications\) on page 1559](#)
- [inactivity-non-tcp-timeout on page 1560](#)
- [inactivity-tcp-timeout on page 1560](#)
- [ip on page 1561](#)
- [local-policy-decision-function on page 1562](#)

- [log \(aACL\) on page 1563](#)
- [match-direction on page 1563](#)
- [max-checked-bytes on page 1564](#)
- [maximum-transactions on page 1564](#)
- [member on page 1565](#)
- [min-checked-bytes on page 1565](#)
- [nested-application on page 1566](#)
- [nested-application-settings on page 1567](#)
- [no-application-identification on page 1567](#)
- [no-application-system-cache on page 1568](#)
- [no-clear-application-system-cache on page 1568](#)
- [no-nested-application on page 1569](#)
- [no-protocol-method on page 1569](#)
- [no-signature-based on page 1570](#)
- [order \(Services Application Identification\) on page 1570](#)
- [pattern on page 1571](#)
- [policy-decision-statistics-profile on page 1572](#)
- [port-mapping on page 1573](#)
- [port-range on page 1573](#)
- [profile on page 1574](#)
- [protocol on page 1574](#)
- [rule \(AACL Rule Set\) on page 1575](#)
- [rule \(Application Identification\) on page 1576](#)
- [rule \(Including in Rule Set\) on page 1577](#)
- [rule-set \(Services AACL\) on page 1577](#)
- [rule-set \(Services Application Identification\) on page 1578](#)
- [service-set-options on page 1578](#)
- [statistics \(System Services\) on page 1579](#)
- [support-uni-directional-traffic on page 1579](#)
- [service-set \(Services\) on page 1580](#)
- [services \(AACL\) on page 1582](#)
- [services \(Application Identification\) on page 1582](#)
- [session-timeout \(Application Identification\) on page 1583](#)
- [session-timeout \(Interfaces\) on page 1583](#)
- [signature on page 1584](#)
- [signature-method-all-ports on page 1584](#)
- [source on page 1585](#)

- [source-address \(AACL\) on page 1585](#)
- [source-address-range on page 1586](#)
- [source-prefix-list \(Services AACL\) on page 1586](#)
- [source-prefix-list \(Services IDS\) on page 1587](#)
- [term on page 1588](#)
- [then on page 1589](#)
- [traceoptions \(Application Identification\) on page 1591](#)
- [traceoptions \(Services Local Policy Decision Function\) on page 1593](#)
- [type on page 1594](#)
- [type-of-service on page 1595](#)
- [url on page 1595](#)

[edit services aacl] Hierarchy List

To configure application-aware access list (AACL) services, include the **aacl** statements at the **[edit services]** hierarchy level:

```
[edit services]
aac1 {
  rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
      from {
        application-group-any;
        application-groups [ application-group-names ];
        application-unknown;
        applications [ application-names ];
        destination-address address <any-unicast>;
        destination-address-range low minimum-value high maximum-value;
        destination-prefix-list list-name;
        source-address address <any-unicast>;
        source-address-range low minimum-value high maximum-value;
        source-prefix-list list-name;
      }
      then {
        (accept | discard);
        count (application | application-group | application-group-any | none);
        forwarding-class class-name;
        policer policer-name;
      }
    }
  }
  rule-set rule-set-name {
    [ rule rule-names ];
  }
}
```

- Related Documentation**
- [Configuring AACL Rules on page 661](#)
 - [Configuring AACL Rule Sets on page 666](#)

- [Configuring Logging of AACL Flows on page 667](#)

[edit services application-identification] Hierarchy Level

To configure application identification services (APPID), include the **application-identification** statement at the [edit services] hierarchy level:

```
[edit services]
application-identification {
  application application-name {
    disable;
    idle-timeout seconds;
    index number;
    session-timeout seconds;
    type type;
    type-of-service service-type;
    port-mapping {
      port-range {
        tcp (port | range);
        udp (port | range);
      }
      disable;
    }
  }
}
application-group group-name {
  application-groups {
    name [application-group-name];
  }
  applications {
    name [application-name];
  }
  index number;
  disable;
}
application-system-cache-timeout seconds;
enable-heuristics
max-checked-bytes bytes;
min-checked-bytes bytes;
nested-application
nested-application-settings
no-application-identification;
no-application-system-cache;
no-clear-application-system-cache;
no-protocol-method;
no-signature-based;
profile profile-name {
  [ rule-set rule-set-name ];
}
rule rule-name {
  disable;
  address address-name {
    destination {
      ip address</prefix-length>;
      port-range {
```

```

        tcp [ ports-and-port-ranges ];
        udp [ ports-and-port-ranges ];
    }
}
source {
    ip address </prefix-length>;
    port-range {
        tcp [ ports-and-port-ranges ];
        udp [ ports-and-port-ranges ];
    }
}
order number;
}
application application-name;
}
rule-set rule-set-name {
    rule application-rule-name;
}
signature-method-all-ports
traceoptions {
    file filename <files number> <match regex> <size size> <world-readable |
        no-world-readable>;
    flag flag;
    no-remote-trace;
}
}
[edit services]
hcm {
    url-rule url-rule-name {
        term term-num {
            from {
                url-list url-list-name ;
                url url_identifier {
                    host hostname ;
                    request-url page-name ;
                }
            }
            then {
                discard;
                accept;
                count;
                log-request;
            }
        }
    }
    url-rule-set url-rule-set-name {
        url-rule rule1 ;
        url-rule rule2 ;
    }
}
}

```

**Related
Documentation**

- [Defining an Application Identification on page 674](#)
- [Application Identification for Nested Applications on page 681](#)

- [Configuring Global APPID Properties on page 683](#)

aacl-fields

| | |
|---------------------------------|--|
| Syntax | <code>aacl-fields {
 <i>field-name</i>;
}</code> |
| Hierarchy Level | [edit system services local-policy-decision-function statistics aacl-statistics-profile <i>profile-name</i>] |
| Release Information | Statement introduced in Junos OS Release 10.0.
IPv6 support introduced in Junos OS Release 12.2 |
| Description | Define the statistics to collect in a data log file. |
| Options | <p><i>field-name</i>—Name of the field:</p> <ul style="list-style-type: none"> • address—IPv4 address • all-fields—All available fields • application—Application name • application-group—Application group name • input-bytes—Number of input bytes • input-interface—Input interface name • input-packets—Number of input packets • ipv6-address—IPv6 address • ipv6-prefix-length—Prefix length associated with the displayed IPv6 address • mask—Netmask • output-bytes—Number of output bytes • output-packets—Number of output packets • subscriber-name—Subscriber name • timestamp—Timestamp • vrf-name—VPN routing and forwarding (VRF) name |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Statistics Profiles on page 702 |

aacl-statistics-profile

| | |
|---------------------------------|---|
| Syntax | <pre>aacl-statistics-profile <i>profile-name</i> {
 aacl-fields {
 <i>field-name</i>;
 }
 file <i>filename</i>;
 record-mode (interim-active-only interim-full);
 report-interval <i>minutes</i>;
}</pre> |
| Hierarchy Level | [edit services service-set <i>service-set-name</i>],
[edit system services local-policy-decision-function statistics] |
| Release Information | Statement introduced in Junos OS Release 10.0.
record-mode option introduced in Junos OS Release 10.2. |
| Description | Create an ACL statistics profile, which configures the files to which statistics records are exported and the format that is exported. |
| Options | <p>file <i>filename</i>—Name of the file to receive the statistics data output. Enclose the name within quotation marks. All files are placed in the directory <code>/var/stats/aacl</code>.</p> <p>record-mode—Record mode for the reporting interval; possible values are interim-active-only, which reports only statistics that have changed, or interim-full, which reports all available statistics.</p> <p>report-interval <i>minutes</i>—Frequency at which statistics are recorded, in minutes.</p> <p>Default: 15 minutes</p> <p>Range: 5 through 1440 minutes</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">For more information on profiles, see the <i>Network Management Administration Guide for Routing Devices</i>.Configuring Statistics Profiles on page 702 |

address

```
Syntax  address address-name {
          destination {
            ip address</prefix-length>;
            port-range {
              tcp [ ports-and-port-ranges ];
              udp [ ports-and-port-ranges ];
            }
          }
          source {
            ip address</prefix-length>;
            port-range {
              tcp [ ports-and-port-ranges ];
              udp [ ports-and-port-ranges ];
            }
          }
          order number;
        }
```

Hierarchy Level [edit services application-identification **rule** *rule-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define address properties for application-identification rule processing. This statement is mandatory; you must specify either the destination or source properties.

Options *address-name*—Identifier for address information.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring APPID Rules on page 676](#)

application (Defining)

Syntax `application application-name {
 disable;
 idle-timeout seconds;
 index number;
 port-mapping {
 disable;
 port-range {
 tcp [ports-and-port-ranges];
 udp [ports-and-port-ranges];
 }
 }
 session-timeout seconds;
 type type;
 type-of-service service-type;
 }`

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define the application and its properties.

The remaining statements are explained separately.

Options ***application-name***—Identifier for the application. This is a mandatory value and has a maximum length of 32 characters.

Required Privilege interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

Related Documentation

- [Defining an Application Identification on page 674](#)

application (Including in Rule)

| | |
|---------------------------------|---|
| Syntax | <code>application <i>application-name</i>;</code> |
| Hierarchy Level | [edit services application-identification rule <i>rule-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Identify the application for inclusion in a rule. |
| Options | <i>application-name</i> —Identifier for the application. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring APPID Rules on page 676 |

application-aware-access-list-fields

| | |
|---------------------------------|---|
| Syntax | <pre>application-aware-access-list-fields {
 <i>field-name</i>;
}</pre> |
| Hierarchy Level | [edit accounting-options policy-decision-statistics-profile <i>profile-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define the statistics to collect in a data log file. |
| Options | <p><i>field-name</i>—Name of the field:</p> <ul style="list-style-type: none">• address—IP address• application—Application name• application-group—Application group name• input-bytes—Number of input bytes• input-interface—Input interface name• input-packets—Number of input packets• mask—Netmask• output-bytes—Number of output bytes• output-packets—Number of output packets• subscriber-name—Subscriber name• timestamp—Timestamp• vrf-name—VPN routing and forwarding (VRF) name |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Statistics Profiles on page 702 |

application-group

| | |
|---------------------------------|--|
| Syntax | <pre> application-group <i>group-name</i> { disable; application-groups { <i>application-group-name</i>; } applications { <i>application-name</i>; } index <i>number</i>; } </pre> |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define the properties and contents of the application group. |
| Options | <p><i>group-name</i>—Unique identifier for the group.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Application Groups on page 680 |

application-group-any

| | |
|---------------------------------|--|
| Syntax | application-group-any; |
| Hierarchy Level | [edit services aacl rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Indicates that any application group defined in the database is considered a match. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Match Conditions in AACL Rules on page 662 |

application-groups (Services ACL)

| | |
|---------------------------------|---|
| Syntax | <code>application-groups [<i>application-group-names</i>];</code> |
| Hierarchy Level | [edit services aacl rule rule-name term term-name from] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Identify one or more application groups defined in the application identification configuration for inclusion as a match condition. |
| Options | <i>application-group-names</i> —Identifiers of the application groups. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Match Conditions in ACL Rules on page 662 |

application-groups (Services Application Identification)

| | |
|---------------------------------|---|
| Syntax | <code>application-groups {
 <i>application-group-name</i>;
}</code> |
| Hierarchy Level | [edit services application-identification application-group group-name] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Identify the list of application groups for inclusion in a larger application group. An <i>application-group-name</i> statement is mandatory. |
| Options | <i>application-group-name</i> —Identifier for the application group. Maximum length is 32 characters. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Application Groups on page 680 |

application-system-cache-timeout

| | |
|---------------------------------|---|
| Syntax | <code>application-system-cache-timeout <i>seconds</i>;</code> |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Configure the lifetime for entries in the application system cache. |
| Options | <i>seconds</i> —Lifetime for system cache entries, in seconds. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Global APPID Properties on page 683 |

application-unknown

| | |
|------------------------------|--|
| Syntax | <code>application-unknown</code> |
| Hierarchy Level | [edit services aacl rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 11.2. |
| Description | Enable AACL logging of flows for unknown applications. |
| Related Documentation | <ul style="list-style-type: none"> • See Configuring Logging of AACL Flows on page 667. |

applications (Services AACL)

| | |
|---------------------------------|---|
| Syntax | <code>applications [<i>application-names</i>];</code> |
| Hierarchy Level | [edit services aacl rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Identify one or more applications defined in the application identification configuration for inclusion as a match condition. |
| Options | <i>application-names</i> —Identifiers of the applications. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Match Conditions in AACL Rules on page 662 |

applications (Services Application Identification)

| | |
|---------------------------------|---|
| Syntax | <pre>applications {
 <i>application-name</i>;
}</pre> |
| Hierarchy Level | [edit services application-identification application-group group-name] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Identify the list of applications for inclusion in the application group. |
| Options | <i>application-name</i> —Identifier for the application. Maximum length is 32 characters. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Application Groups on page 680 |

automatic

| | |
|---------------------------------|---|
| Syntax | <pre>automatic {
 interval <i>hour</i>;
 start-time <i>time</i>;
}</pre> |
| Hierarchy Level | [edit services application-identification download] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define automatic download properties. |
| Options | <i>interval hour</i> —Download interval in hours. The default is 24 and the range is from 1 through 168.

<i>start-time time</i> —Start-time value. The default is 0:00 and the range is from 0:00 through 24:00. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Automatic Download of Application Package Updates on page 686 |

bypass-traffic-on-exceeding-flow-limits

| | |
|---------------------------------|---|
| Syntax | bypass-traffic-on-exceeding-flow-limits; |
| Hierarchy Level | [edit services service-set <i>service-set-name</i> service-set-options] |
| Release Information | Statement introduced in Junos OS Release 10.1. |
| Description | Enable packets to bypass without creating a new session when the flow in the service set exceeds the limit that is set by the max-flows statement at the [edit services service-set <i>service-set-name</i>] hierarchy level. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Service Sets to be Applied to Services Interfaces on page 31 |

chain-order

| | |
|---------------------------------|--|
| Syntax | chain-order; |
| Hierarchy Level | [edit services application-identification nested-application <i>name</i> signature <i>name</i>] |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Signatures can contain multiple members. If the chain order feature is on, those members are read in order. By default, chain ordering is turned off. If a signature contains only one member, this option is ignored. |
| Required Privilege Level | system—To view this statement in the configuration.
system control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Application Identification for Nested Applications on page 681 |

context

| | |
|---------------------------------|--|
| Syntax | context (http-header-content-type http-header-host http-url-parsed http-url-parsed-param-parsed); |
| Hierarchy Level | [edit services application-identification nested-application name signature name member name] |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Define a service-specific context, such as http-url . |
| Options | value —Service-specific context. |
| Required Privilege Level | system—To view this statement in the configuration.
system control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Application Identification for Nested Applications on page 681 |

destination (Services)

| | |
|---------------------------------|---|
| Syntax | <pre>destination {
 ip address</prefix-length>;
 port-range {
 tcp [<i>ports-and-port-ranges</i>];
 udp [<i>ports-and-port-ranges</i>];
 }
}</pre> |
| Hierarchy Level | [edit services application-identification rule rule-name address address-name] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define destination properties for application-identification rule processing. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring APPID Rules on page 676 |

destination-address

| | |
|---------------------------------|--|
| Syntax | <code>destination-address <i>address</i>;</code> |
| Hierarchy Level | [edit services aacl rule rule-name term term-name from] |
| Release Information | Statement introduced in Junos OS Release 9.5.
IPv6 support introduced in Junos OS Release 12.2 |
| Description | Specify the destination address for rule matching. |
| Options | <i>address</i> —Destination IPv4 or IPv6 address or prefix value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Match Conditions in AACL Rules on page 662 |

destination-address-range

| | |
|---------------------------------|--|
| Syntax | <code>destination-address-range low <i>minimum-value</i> high <i>maximum-value</i>;</code> |
| Hierarchy Level | [edit services aacl rule rule-name term term-name from] |
| Release Information | Statement introduced in Junos OS Release 9.5.
IPv6 support introduced in Junos OS Release 12.2 |
| Description | Specify the destination address range for rule matching. |
| Options | <i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range.
<i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Match Conditions in AACL Rules on page 662 |

destination-prefix-list

| | |
|--------------------------|---|
| Syntax | <code>destination-prefix-list <i>list-name</i>;</code> |
| Hierarchy Level | <code>[edit services aacl rule <i>rule-name</i> term <i>term-name</i> from]</code> |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Specify the destination prefix list for rule matching. You configure the prefix list by including the prefix-list statement at the <code>[edit policy-options]</code> hierarchy level. |
| Options | <i>list-name</i> —Destination prefix list. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Match Conditions in ACL Rules on page 662 |

direction

| | |
|--------------------------|--|
| Syntax | <code>direction (any client-to-server server-to-client) ;</code> |
| Hierarchy Level | <code>[edit services application-identification nested-application <i>name</i> signature <i>name</i> member <i>name</i>]</code> |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Specify the connection direction of the packets to apply pattern matching. |
| Options | <i>direction</i> —The directions of packets are client-to-server , server-to-client , or any . |
| Required Privilege Level | system—To view this statement in the configuration.
system control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Application Identification for Nested Applications on page 681. |

disable (APPID Application)

| | |
|---------------------------------|---|
| Syntax | disable; |
| Hierarchy Level | [edit services application-identification application <i>application-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Disable this application definition. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Defining an Application Identification on page 674 |

disable (APPID Application Group)

| | |
|---------------------------------|---|
| Syntax | disable; |
| Hierarchy Level | [edit services application-identification application-group <i>group-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Disable application group properties. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Application Groups on page 680 |

disable (APPID Port Mapping)

| | |
|---------------------------------|---|
| Syntax | disable; |
| Hierarchy Level | [edit services application-identification application <i>application-name</i> port-mapping] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Disable port-mapping properties for application identification. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Defining an Application Identification on page 674 |

disable-global-timeout-override

| | |
|---------------------------------|---|
| Syntax | disable-global-timeout-override; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options] |
| Release Information | Statement introduced in Junos OS Release 10.0. |
| Description | Disallow overriding a global inactivity or session timeout. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Defining an Application Identification on page 674 |

download

| | |
|---------------------------------|---|
| Syntax | <pre>download {
 automatic {
 interval <i>hour</i>;
 start-time <i>time</i>;
 }
 url <i>url</i>;
}</pre> |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define application download properties. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Automatic Download of Application Package Updates on page 686 |

enable-asymmetric-traffic-processing

| | |
|---------------------------------|--|
| Syntax | enable-asymmetric-traffic-processing; |
| Hierarchy Level | [edit services service-set <i>service-set-name</i> service-set-options] |
| Release Information | Statement introduced in Junos OS Release 11.2. |
| Description | Enables APPID to perform application matching on unidirectional traffic. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring APPID Support for Unidirectional Traffic on page 685 |

enable-heuristics

| | |
|---------------------------------|--|
| Syntax | enable-heuristics; |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 11.2. |
| Description | Enables APPID to identify encrypted data packets in point-to-point applications by using heuristics methodology. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring APPID Support for Heuristics on page 684 |

file

| | |
|---------------------------------|---|
| Syntax | <pre>file <i>file-name</i> {
 archive-sites <i>url</i>;
 files <i>file-number</i>;
 size <i>bytes</i>;
 transfer-interval <i>minutes</i>;
}</pre> |
| Hierarchy Level | [edit system services local-policy-decision-function statistics] |
| Release Information | Statement introduced in Junos OS Release 10.0. |
| Description | Specify a file to which statistics records are exported and the format that is exported. |
| Options | <p>archive-sites [<i>url</i>]—One or more destinations for archiving data.</p> <p>filename—Name of the file to receive the statistics data output.</p> <p>files <i>number</i>—(Optional) Maximum number of accounting files.
Range: 3 through 1000 files
Default: 3 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>size <i>size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).
Syntax: <i>xk</i> to specify KB, <i>xm</i> to specify MB, or <i>xg</i> to specify GB
Range: 262144 through 1073741824 or the maximum file size supported on your system</p> <p>If you specify a maximum file size, you also must specify a maximum number of trace files with the files option.</p> <p>transfer-interval <i>minutes</i>—Frequency at which to transfer files to archive sites, in minutes.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Statistics Profiles on page 702 |


from

| | |
|---------------------------------|--|
| Syntax | <pre> from { application-group-any; application-groups [application-group-names]; application-unknown; applications [application-names]; destination-address address <any-unicast>; destination-address-range low minimum-value high maximum-value; destination-prefix-list list-name; nested-application-unknown; source-address address <any-unicast>; source-address-range low minimum-value high maximum-value; source-prefix-list list-name; } </pre> |
| Hierarchy Level | [edit services aacl rule rule-name term term-name] |
| Release Information | Statement introduced before Junos OS Release 9.5. |
| Description | Specify match conditions for the AACL term. |
| Options | <p>For information on match conditions, see the description of firewall filter match conditions in the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring AACL Rules on page 661 |

idle-timeout

| | |
|---------------------------------|---|
| Syntax | <code>idle-timeout seconds;</code> |
| Hierarchy Level | [edit services application-identification application <i>application-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define idle timeout for an application in seconds. When the timeout period expires, the session ends if no packets have been received. |
| Options | seconds —Idle timeout period.
Default: 30
Range: 1 through 604,800 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• APPID Overview on page 669• Defining an Application Identification on page 674 |

ignore-errors

| | |
|---|--|
| Syntax | <code>ignore-errors <alg> <tcp>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options] |
| Release Information | Statement introduced in Junos OS Release 10.1. |
| Description | Define settings for minimizing TCP packet drops during stateful firewall processing. |
| <div> NOTE: <code>ignore-errors</code> option is not supported on adaptive services interfaces (<code>sp-x/y/z</code>).</div> | |
| Options | alg —Mediate ALG behavior that results in dropping malformed packets or random packets when the software is unable to allocate resources.
tcp —Prevent software from dropping packets that fail TCP SYN checks. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Defining an Application Identification on page 674 |

index (Applications)

| | |
|---------------------------------|--|
| Syntax | <code>index number;</code> |
| Hierarchy Level | [edit services application-identification application application-name],
[edit services application-identification application-group group-name] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Assign an application or application-group index number. This is a mandatory value. |
| Options | number —Index number; must be a unique, unsigned value.
Range: 0 through 65535 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Defining an Application Identification on page 674 • Configuring Application Groups on page 680 |

index (Nested Applications)

| | |
|---------------------------------|--|
| Syntax | <code>index number;</code> |
| Hierarchy Level | [edit services application-identification nested-application name] |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Set a number that is a one-to-one mapping to the application name. The application name is used to ensure that each signature definition is unique. |
| Options | number —Numeric value associated with an application name. The index range for predefined applications is from 1 through 32767. The index range for custom applications and custom nested applications is from 32768 through 65534. |
| Required Privilege Level | system—To view this statement in the configuration.
system control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Application Identification for Nested Applications on page 681. |

inactivity-non-tcp-timeout

| | |
|---------------------------------|---|
| Syntax | <code>inactivity-non-tcp-timeout <i>seconds</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options] |
| Release Information | Statement introduced in Junos OS Release 10.0. |
| Description | Define the inactivity timeout period for non-TCP established sessions in seconds. |
| Options | <i>seconds</i> —Timeout period.
Range: 4 through 86,400 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Defining an Application Identification on page 674 |

inactivity-tcp-timeout

| | |
|---------------------------------|---|
| Syntax | <code>inactivity-tcp-timeout <i>seconds</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options] |
| Release Information | Statement introduced in Junos OS Release 10.0. |
| Description | Define the inactivity timeout period for TCP established sessions in seconds. |
| Options | <i>seconds</i> —Timeout period.
Range: 4 through 86,400 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Defining an Application Identification on page 674 |

ip

| | |
|---------------------------------|--|
| Syntax | <code>ip address</prefix-length>;</code> |
| Hierarchy Level | [edit services application-identification rule rule-name address destination],
[edit services application-identification rule rule-name address source] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define an IP address and netmask for identifying the traffic destination or source. |
| Options | <code>address</prefix-length></code> —IP address and netmask. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring APPID Rules on page 676 |

local-policy-decision-function

Syntax local-policy-decision-function {
 statistics {
 aocl-statistics-profile *profile-name* {
 aocl-fields {
 field-name;
 }
 file *filename*;
 report-interval *minutes*;
 }
 file *file-name* {
 archive-sites *url*;
 files *file-number*;
 size *bytes*;
 transfer-interval *minutes*;
 }
 record-type (delta | interim);
 }
 traceoptions {
 file *filename* <files *number*> <size *size*>;
 flag *flag*;
 no-remote-trace;
 }
 }

Hierarchy Level [edit system services]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify L-PDF properties.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Configuring Statistics Profiles on page 702](#)

log (aACL)

| | |
|---------------------------------|--|
| Syntax | log <i>event-type</i> |
| Hierarchy Level | [edit services aACL rule <i>rule-name</i> term <i>term-name</i> then] |
| Release Information | Statement introduced in Junos OS Release 11.2. |
| Description | Enable AACL logging of flows for known or unknown applications. |
| Options | <i>event-type</i> —Enable logging of the specified <i>event-type</i> : <ul style="list-style-type: none"> • session-start • session-end • session-start-end-no-stats • session-start-interim-end • session-interim end • session-end |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • See Configuring Logging of AACL Flows on page 667. |

match-direction

| | |
|---------------------------------|--|
| Syntax | match-direction (input output input-output); |
| Hierarchy Level | [edit services aACL rule <i>rule-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Specify the direction in which the rule match is applied. |
| Options | <i>input</i> —Apply the rule match on the input side of the interface.

<i>output</i> —Apply the rule match on the output side of the interface.

<i>input-output</i> —Apply the rule match bidirectionally. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Match Direction for AACL Rules on page 662 |

max-checked-bytes

| | |
|---------------------------------|---|
| Syntax | <code>max-checked-bytes <i>bytes</i>;</code> |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Specify the maximum number of bytes to be inspected. |
| Options | <i>bytes</i> —Maximum number of bytes.
Range: 0 through 100,000 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Global APPID Properties on page 683 |

maximum-transactions

| | |
|---------------------------------|--|
| Syntax | <code>maximum-transactions <i>number</i>;</code> |
| Hierarchy Level | [edit services application-identification <i>nested-application name signature name</i>] |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Set the maximum number of transactions required before a match is made. |
| Options | <i>number</i> —Maximum number of transactions. |
| Required Privilege Level | system—To view this statement in the configuration.
system control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Application Identification for Nested Applications on page 681 |

member

| | |
|---------------------------------|---|
| Syntax | <code>member <i>name</i>;</code> |
| Hierarchy Level | [edit services application-identification nested-application <i>name</i> signature <i>name</i>] |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Define a member name for a custom nested application signature definition. Custom definitions can contain multiple members that define attributes for an application. |
| Options | <i>name</i> —Name of member for a custom nested application signature definition. |
| Required Privilege Level | system—To view this statement in the configuration.
system control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Application Identification for Nested Applications on page 681 |

min-checked-bytes

| | |
|---------------------------------|---|
| Syntax | <code>min-checked-bytes <i>bytes</i>;</code> |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Specify the minimum number of bytes to be inspected. |
| Options | <i>bytes</i> —Minimum number of bytes.
Range: 0 through 2000 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Global APPID Properties on page 683 |

nested-application

| | |
|--------------------------|---|
| Syntax | <pre>nested-application <i>name</i> {
 <i>index number</i>;
 <i>protocol protocol</i> ;
 <i>signature name</i> {
 <i>chain-order</i> ;
 <i>maximum-transactions number</i>;
 <i>member name</i> {
 <i>context</i> (http-header-content-type http-header-host http-url-parsed
 http-url-parsed-param-parsed);
 <i>direction</i> (any client-to-server server-to-client);
 <i>pattern dfa-pattern</i>;
 }
 <i>order number</i>;
 }
 <i>type type</i>;
}</pre> |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Configure a custom nested application definition for the desired application name that will be used by the system to identify the nested application as it passes through the device. Custom nested application definitions can be used for nested applications that are not part of the Juniper Networks predefined nested application database. |
| Options | <p><i>name</i>—Name of nested application.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Application Identification for Nested Applications on page 681 |

nested-application-settings

| | |
|---------------------------------|---|
| Syntax | nested-application-settings {
no-application-system-cache;
no-nested-application;
} |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Configure nested application options for application identification services. |
| Required Privilege Level | system—To view this statement in the configuration.
system control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Application Identification for Nested Applications on page 681. |

no-application-identification

| | |
|---------------------------------|---|
| Syntax | no-application-identification; |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Disable all application identification methods. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Global APPID Properties on page 683 |

no-application-system-cache

| | |
|---------------------------------|--|
| Syntax | no-application-system-cache; |
| Hierarchy Level | [edit services application-identification],
[edit services application-identification nested-application-settings] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Disable storing application identification results in the application system cache. Nested application identification information is saved in the application system cache to improve performance. This cache is updated when a different application is identified. This caching is turned on by default. Use the no-application-system-cache statement to turn it off. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Global APPID Properties on page 683• Application Identification for Nested Applications on page 681. |

no-clear-application-system-cache

| | |
|---------------------------------|---|
| Syntax | no-clear-application-system-cache; |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Disable clearing the application system cache. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Global APPID Properties on page 683 |

no-nested-application

| | |
|---------------------------------|--|
| Syntax | no-nested-application; |
| Hierarchy Level | [edit services application-identification nested-application-settings] |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Sometimes there is a need to identify multiple different applications running on the same Layer 7 protocols. For example, both Facebook and Yahoo Messenger can run over HTTP, but there is a need to identify them as two different applications. To do this, the application layer is split into two layers: Layer 7 applications and Layer 7 protocols. This function is turned on by default. Use the no-nested-application statement to turn it off. |
| Required Privilege Level | system—To view this statement in the configuration.
system control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Application Identification for Nested Applications on page 681 |

no-protocol-method

| | |
|---------------------------------|---|
| Syntax | no-protocol-method; |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 10.1. |
| Description | Disable the protocol-based application identification method. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Global APPID Properties on page 683 |

no-signature-based

| | |
|---------------------------------|---|
| Syntax | no-signature-based; |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Disable the signature-based application identification method. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Global APPID Properties on page 683 |

order (Services Application Identification)

| | |
|---------------------------------|--|
| Syntax | order <i>number</i> ; |
| Hierarchy Level | [edit services application-identification <i>nested-application name signature name member name</i>]
[edit services application-identification <i>rule rule-name address</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define application matching priority. For address configurations, the order number resolves the conflict when multiple address entries are matched for a specific session. The lower number has higher priority. |
| Options | <i>number</i> —Order number. This value is mandatory and must be unique. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring APPID Rules on page 676• Application Identification for Nested Applications on page 681 |

pattern

| | |
|---------------------------------|--|
| Syntax | <code>pattern <i>dfa-pattern</i>;</code> |
| Hierarchy Level | [edit services application-identification nested-application name signature name member name] |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Define an attack pattern to be detected. |
| Options | <i>dfa-pattern</i> —Pattern of attack to match. Deterministic Finite Automata (DFA) is a powerful pattern matching engine. |
| Required Privilege Level | system—To view this statement in the configuration.
system control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Application Identification for Nested Applications on page 681 |

policy-decision-statistics-profile

| | |
|---------------------------------|---|
| Syntax | <pre>policy-decision-statistics-profile <i>profile-name</i> {
 aacl-fields {
 <i>field-name</i>;
 }
 file <i>filename</i>;
 files <i>file-number</i>;
 size <i>bytes</i>;
}</pre> |
| Hierarchy Level | [edit accounting-options],
[edit services service-set <i>service-set-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Create a policy decision statistics profile, which configures the files to which statistics records are exported and the format that is exported. |
| Options | <p>file <i>filename</i>—Name of the file to receive the accounting-data output. Enclose the name within quotation marks. All files are placed in the directory /var/log.</p> <p>files <i>number</i>—(Optional) Maximum number of accounting files.
Range: 2 through 1000 files
Default: 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p><i>profile-name</i>—Name of the policy decision statistics profile.</p> <p>size <i>size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).
Syntax: xk to specify KB, xm to specify MB, or xg to specify GB
Range: 10240 through 1073741824 or the maximum file size supported on your system</p> <p>If you specify a maximum file size, you also must specify a maximum number of trace files with the files option.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">For more information on profiles, see the <i>Network Management Administration Guide for Routing Devices</i>.Configuring Statistics Profiles on page 702 |

port-mapping

| | |
|---------------------------------|---|
| Syntax | <pre>port-mapping { disable; port-range { tcp [ports-and-port-ranges]; udp [ports-and-port-ranges]; } }</pre> |
| Hierarchy Level | [edit services application-identification application <i>application-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define port-mapping properties for application identification. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Defining an Application Identification on page 674 |

port-range

| | |
|---------------------------------|---|
| Syntax | <pre>port-range { tcp [ports-and-port-ranges]; udp [ports-and-port-ranges]; }</pre> |
| Hierarchy Level | [edit services application-identification application <i>application-name</i> port-mapping],
[edit services application-identification rule <i>rule-name</i> address destination],
[edit services application-identification rule <i>rule-name</i> address source] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define TCP and UDP port numbers or numeric ranges. For port-mapping configurations, this entry is required if the parent node exists. |
| Options | ports-and-port-ranges —Individual port numbers, numeric port ranges, or both. Separate the values with spaces. The format for numeric port ranges is <i>minimum-value–maximum-value</i> . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Defining an Application Identification on page 674 • Configuring APPID Rules on page 676 |

profile

| | |
|---------------------------------|---|
| Syntax | <code>profile <i>profile-name</i> {
 rule-set <i>rule-set-name</i>;
}</code> |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define members of application profile, which consists of one or more rule sets. |
| Options | <i>profile-name</i> —Identifier for application profile.

The remaining statement is explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Application Profiles on page 679 |

protocol

| | |
|---------------------------------|--|
| Syntax | <code>protocol <i>protocol</i>;</code> |
| Hierarchy Level | [edit services application-identification nested-application <i>name</i>] |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Identify the protocol that will be monitored to identify nested applications. HTTP is supported. |
| Options | <i>protocol</i> —An agreed-upon or standardized method for transmitting data and establishing communications between different devices. The value http is supported. |
| Required Privilege Level | system—To view this statement in the configuration.
system control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Application Identification for Nested Applications on page 681 |

rule (AACL Rule Set)

| | |
|---------------------------------|--|
| Syntax | <pre> rule <i>rule-name</i> { match-direction (input output input-output); term <i>term-name</i> { from { application-group-any; application-groups [<i>application-group-names</i>]; application-unknown; applications [<i>application-names</i>]; destination-address <i>address</i> <any-unicast>; destination-address-range low <i>minimum-value</i> high <i>maximum-value</i>; destination-prefix-list <i>list-name</i>; nested-application-unknown; source-address <i>address</i> <any-unicast>; source-address-range low <i>minimum-value</i> high <i>maximum-value</i>; source-prefix-list <i>list-name</i>; } then { (accept discard); count (application application-group application-group-any nested-application none); forwarding-class <i>class-name</i>; policer <i>policer-name</i>; } } } </pre> |
| Hierarchy Level | [edit services aacl],
[edit services aacl rule-set <i>rule-set-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Specify the rule the router uses when applying this service. |
| Options | <p>rule-name—Identifier for the collection of terms that constitute this rule.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring AACL Rules on page 661 |

rule (Application Identification)

Syntax

```
rule rule-name {  
  address {  
    destination {  
      ip address </prefix-length>;  
      port-range {  
        tcp [ ports-and-port-ranges ];  
        udp [ ports-and-port-ranges ];  
      }  
    }  
  }  
  source {  
    ip address </prefix-length>;  
    port-range {  
      tcp [ ports-and-port-ranges ];  
      udp [ ports-and-port-ranges ];  
    }  
  }  
  order number;  
}  
application application-name;  
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define properties for application-identification rule processing.

Options *rule-name*—Unique identifier for the rule.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring APPID Rules on page 676](#)

rule (Including in Rule Set)

| | |
|---------------------------------|---|
| Syntax | <code>rule <i>rule-name</i>;</code> |
| Hierarchy Level | [edit services application-identification rule-set <i>rule-set-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Identify rules for inclusion in application rule set. |
| Options | <i>rule-name</i> —Unique identifier for the rule. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring APPID Rules on page 676 |

rule-set (Services ACL)

| | |
|---------------------------------|---|
| Syntax | <code>rule-set <i>rule-set-name</i> {
 [rule <i>rule-names</i>];
}</code> |
| Hierarchy Level | [edit services aacl] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Specify the rule set the router uses when applying this service. |
| Options | <i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring AACL Rule Sets on page 666 |

rule-set (Services Application Identification)

| | |
|--------------------------|---|
| Syntax | <pre>rule-set <i>rule-set-name</i> {
 <i>rule application-rule-name</i>;
}</pre> |
| Hierarchy Level | [edit services application-identification],
[edit services application-identification <i>profile profile-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define members of rule set. |
| Options | <i>rule-set-name</i> —Unique identifier for the rule set.

The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring APPID Rules on page 676 |

service-set-options

| | |
|--------------------------|---|
| Syntax | <pre>service-set-options {
 <i>bypass-traffic-on-exceeding-flow-limits</i>;
 <i>bypass-traffic-on-pic-failure</i>;
 <i>enable-asymmetric-traffic-processing</i>;
 header-integrity-check
 routing-engine-services;
 <i>support-uni-directional-traffic</i>;
}</pre> |
| Hierarchy Level | [edit services service-set] |
| Release Information | Statement introduced in Junos OS Release 10.1. The enable-asymmetric-traffic-processing and the support-uni-directional-traffic options were added in Junos OS Release 11.2. The routing-engine-services option was added in Junos OS Release 15.1. |
| Description | Specify the service set options to apply to a service set. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Service Sets to be Applied to Services Interfaces on page 31• Configuring APPID Support for Unidirectional Traffic on page 685 |

statistics (System Services)

| | |
|---------------------------------|--|
| Syntax | <pre> statistics { aacl-statistics-profile <i>profile-name</i> { aacl-fields { <i>field-name</i>; } file <i>filename</i>; report-interval <i>minutes</i>; } file <i>file-name</i> { archive-sites [<i>url</i>]; files <i>file-number</i>; size <i>bytes</i>; transfer-interval <i>minutes</i>; } record-type (delta interim); } </pre> |
| Hierarchy Level | [edit system services local-policy-decision-function] |
| Release Information | Statement introduced in Junos OS Release 10.0. |
| Description | Configure file and data specifications for recording AACL statistics. |
| Options | <p>record-type—Record type; possible values are delta or interim.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Statistics Profiles on page 702 |

support-uni-directional-traffic

| | |
|---------------------------------|--|
| Syntax | support-uni-directional-traffic; |
| Hierarchy Level | [edit services service-set <i>service-set-name</i> service-set-options] |
| Release Information | Statement introduced in Junos OS Release 11.2. |
| Description | Enables APPID to perform application matching on unidirectional traffic. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring APPID Support for Unidirectional Traffic on page 685 |

service-set (Services)

```
Syntax  service-set service-set-name {
        allow-multicast;
        extension-service service-name {
            provider-specific-rules-configuration;
        }
        (ids-rules rule-name | ids-rule-sets rule-set-name);
        interface-service {
            service-interface interface-name;
        }
        ipsec-vpn-options {
            anti-replay-window-size bits;
            clear-dont-fragment-bit;
            ike-access-profile profile-name;
            local-gateway address;
            no-anti-replay;
            passive-mode-tunneling;
            trusted-ca [ ca-profile-names ];
            tunnel-mtu bytes;
        }
        ip-reassembly-rules rule-name;
        (ipsec-vpn-rules rule-name | ipsec-vpn-rule-sets rule-set-name);
        max-flows number;
        max-drop-flows {
            ingress ingress-flows;
            egress egress-flows;
        }
        nat-options {
            land-attack-check (ip-only | ip-port);
            max-sessions-per-subscriber session-number;
            stateful-nat64 {
                clear-dont-fragment-bit;
            }
        }
        (nat-rules rule-name | nat-rule-sets rule-set-name);
        next-hop-service {
            inside-service-interface interface-name.unit-number;
            outside-service-interface interface-name.unit-number;
            outside-service-interface-type local;
            service-interface-pool name;
        }
        (pgcp-rules rule-name | pgcp-rule-sets rule-set-name);
        (ptsp-rules rule-name | ptsp-rule-sets rule-set-name);
        service-set-options {
            bypass-traffic-on-exceeding-flow-limits;
            bypass-traffic-on-pic-failure;
            enable-asymmetric-traffic-processing;
            routing-engine-services;
            support-uni-directional-traffic;
        }
        snmp-trap-thresholds {
            flows high high-threshold | low low-threshold;
            nat-address-port high-threshold | low low-threshold;
        }
    }
```

```

    }
  }
  software-options {
    dslite-ipv6-prefix-length dslite-ipv6-prefix-length;
  }
  (software-rules rule-name | software-rule-sets rule-set-name);
  (stateful-firewall-rules rule-name | stateful-firewall-rule-sets rule-set-name);
  syslog {
    host hostname {
      class {
        alg-logs;
        ids-logs;
        nat-logs;
        packet-logs;
        pcp-logs;
        session-logs <open | close>;
        stateful-firewall-logs ;
      }
      services severity-level;
      facility-override facility-name;
      interface-service prefix-value;
    }
  }
}

```

Hierarchy Level [edit services]

Release Information Statement introduced before Junos OS Release 7.4.
pgcp-rules and **pgcp-rule-sets** options added in Junos OS Release 8.4.
server-set-options option added in Junos OS Release 10.1.
ptsp-rules and **ptsp-rule-sets** options added in Junos OS Release 10.2.
software-rules and **clear-rule-sets** options added in Junos OS Release 10.4.
software-options option added in Junos OS Release 14.1.

Description Define the service set.

Options *service-set-name*—Name of the service set.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- *Service Set Properties*

services (AACL)

| | |
|---------------------------------|---|
| Syntax | <code>services aacl { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | aacl statement introduced in Junos OS Release 9.5. |
| Description | Define the services to be applied to traffic. |
| Options | aacl —The values configured for application-aware-access-list matching rules.

The statements are explained separately. |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Application Aware Services Interfaces Feature Guide for Routing Devices</i> |

services (Application Identification)

| | |
|---------------------------------|--|
| Syntax | <code>services application-identification { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | services statement introduced before Junos OS Release 7.4.
application-identification statement introduced in Junos OS Release 9.5. |
| Description | Define the services to be applied to traffic. |
| Options | application-identification —The values configured for application-identification properties.

The statements are explained separately. |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Application Identification</i> |

session-timeout (Application Identification)

| | |
|---------------------------------|---|
| Syntax | <code>session-timeout <i>seconds</i>;</code> |
| Hierarchy Level | [edit services application-identification application <i>application-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define session lifetime for the specified application in seconds. |
| Options | <i>seconds</i> —Duration of session.
Default: 3600
Range: 1 through 604,800 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Defining an Application Identification on page 674 |

session-timeout (Interfaces)

| | |
|---------------------------------|---|
| Syntax | <code>session-timeout <i>seconds</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options] |
| Release Information | Statement introduced in Junos OS Release 10.0. |
| Description | Define session lifetime globally for the Multiservices interface in seconds. |
| Options | <i>seconds</i> —Duration of session.
Range: 4 through 86,400 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Defining an Application Identification on page 674 |

signature

| | |
|---------------------------------|--|
| Syntax | <pre>signature <i>name</i> {
 chain-order;
 maximum-transactions <i>number</i>;
 member <i>name</i> {
 context <i>value</i>;
 direction (any client-to-server server-to-client);
 pattern <i>dfa-pattern</i>;
 }
 order <i>number</i>;
}</pre> |
| Hierarchy Level | [edit services application-identification nested-application name] |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Identify the name of the custom nested application signature definition. The name must be unique with a maximum length of 32 characters. |
| Options | <p><i>name</i>—Name of the signature definition.</p> <p>The remaining statements are described separately.</p> |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Application Identification for Nested Applications on page 681 |

signature-method-all-ports

| | |
|---------------------------------|--|
| Syntax | <pre>signature-method-all-ports</pre> |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 11.2. |
| Description | <p>Runs signature matching on all traffic in application-identification. This is called the signature-match mode.</p> <p>In the default mode, or fast-port-match mode, all traffic destined to well-known ports (up to 1024) immediately returns the final port match. However, the device runs signature matching for all traffic destined for port 80,</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Global APPID Properties on page 683 |

source

| | |
|---------------------------------|---|
| Syntax | <pre>source { ip address </prefix-length>; port-range { tcp [ports-and-port-ranges]; udp [ports-and-port-ranges]; } }</pre> |
| Hierarchy Level | [edit services application-identification rule <i>rule-name</i> address <i>address-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define source properties for application-identification rule processing. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring APPID Rules on page 676 |

source-address (AACL)

| | |
|---------------------------------|--|
| Syntax | source-address <i>address</i> ; |
| Hierarchy Level | [edit services aacl rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 9.5.
IPv6 support introduced in Junos OS Release 12.2 |
| Description | Specify the source address for rule matching. |
| Options | address —Source IPv4 or IPv6 address or prefix value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Match Conditions in AACL Rules on page 662 |

source-address-range

| | |
|---------------------------------|--|
| Syntax | source-address-range low <i>minimum-value</i> high <i>maximum-value</i> ; |
| Hierarchy Level | [edit services aacl rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 9.5.
IPv6 support introduced in Junos OS Release 12.2 |
| Description | Specify the source address range for rule matching. |
| Options | <i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range.
<i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Match Conditions in AACL Rules on page 662 |

source-prefix-list (Services AACL)

| | |
|---------------------------------|---|
| Syntax | source-prefix-list <i>list-name</i> ; |
| Hierarchy Level | [edit services aacl rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Specify the source prefix list for rule matching. You configure the prefix list by including the prefix-list statement at the [edit policy-options] hierarchy level. |
| Options | <i>list-name</i> —Source prefix list. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Match Conditions in AACL Rules on page 662 |

source-prefix-list (Services IDS)

| | |
|---------------------------------|---|
| Syntax | <code>source-prefix-list <i>list-name</i> <except>;</code> |
| Hierarchy Level | [edit services ids rule <i>rule-name</i> term <i>term-name</i> from] |
| Release Information | Statement introduced in Junos OS Release 8.2. |
| Description | Specify the source prefix list for rule matching. You configure the prefix list by including the prefix-list statement at the [edit policy-options] hierarchy level. |
| Options | <p><i>list-name</i>—Destination prefix list.</p> <p>except—(Optional) Exclude the specified prefix list from rule matching.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Match Conditions in IDS Rules on page 357• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> |

term

Syntax `term term-name {`
 `from {`
 `application-group-any;`
 `application-groups [application-group-names];`
 `application-unknown;`
 `applications [application-names];`
 `destination-address address <any-unicast>;`
 `destination-address-range low minimum-value high maximum-value;`
 `destination-prefix-list list-name;`
 `nested-application-unknown;`
 `source-address address <any-unicast>;`
 `source-address-range low minimum-value high maximum-value;`
 `source-prefix-list list-name;`
 `}`
 `then {`
 `(accept | discard);`
 `count (application | application-group | application-group-any | nested-application |`
 `none);`
 `forwarding-class class-name;`
 `policer policer-name;`
 `}`
 `}`

Hierarchy Level [edit services aacl [rule](#) *rule-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define the AACL term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Configuring AACL Rules on page 661](#)

then

| | |
|----------------------------|--|
| Syntax | <pre> then { (accept discard); count (application application-group application-group-any nested-application none); forwarding-class <i>class-name</i>; log <i>event-type</i>; policer <i>policer-name</i>; } </pre> |
| Hierarchy Level | [edit services aacl rule <i>rule-name</i> term <i>term-name</i>] |
| Release Information | <p>Statement introduced in Junos OS Release 9.5.</p> <p>policer statement added in Junos OS Release 9.6.</p> <p>The nested-application option for the count statement introduced in Junos OS Release 11.1.</p> |
| Description | Define the AACL term actions. You can configure the router to accept or discard the targeted traffic. The action modifiers (count and forwarding-class) are optional. |
| Options | <p>You can configure one of the following actions:</p> <ul style="list-style-type: none"> • accept—Accept the packets and all subsequent packets in flows that match the rules. • discard—Discard the packet and all subsequent packets in flows that match the rules. <p>When you select accept as the action, you can optionally configure one or both of the following action modifiers. No action modifiers are allowed with the discard action.</p> <ul style="list-style-type: none"> • count (application application-group application-group-any nested-application none)—For all accepted packets that match the rules, record a packet count using AACL statistics practices. You can specify one of the following options; there is no default setting: <ul style="list-style-type: none"> • application—Count the application that matched in the from clause. • application-group—Count the application group that matched in the from clause. • application-group-any—Count all application groups that match from application-group-any under the any group name. • nested-application—Count all nested applications that matched in the from clause. • none—Same as not specifying count as an action. • forwarding-class <i>class-name</i>—Specify the packets' forwarding-class name. <p>policer <i>policer-name</i>—Apply rate-limiting properties to the traffic as configured at the [edit firewall policer <i>policer-name</i>] hierarchy level. This configuration allows bit-rate and burst-size attributes to be applied to the traffic that are not supported by AACL rules. When you include a policer, the only allowed action is discard. For more information on policers, see the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>.</p> |

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related • [Configuring ACL Rules on page 661](#)
Documentation • *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*

traceoptions (Application Identification)

| | |
|----------------------------|--|
| Syntax | <pre> traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regex</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; no-remote-trace; } </pre> |
| Hierarchy Level | [edit services application-identification] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | <p>Configure application identification tracing options.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p> |
| Options | <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>, then <i>trace-file.1</i>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>flag—Tracing operation to perform. all is the only valid completion.</p> <ul style="list-style-type: none"> all—Trace all events. <p>match <i>regex</i>—(Optional) Regular expression for lines to be logged.</p> <p>no-world-readable—(Optional) Disallow any user to read the log file.</p> <p>size <i>size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named <i>trace-file</i> reaches this size, it is renamed <i>trace-file.0</i>. When the <i>trace-file</i> again reaches its maximum size, <i>trace-file.0</i> is renamed <i>trace-file.1</i> and <i>trace-file</i> is renamed <i>trace-file.0</i>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Syntax: xk to specify KB, xm to specify MB, or xg to specify GB</p> <p>Range: 10240 through 1073741824 or the maximum file size supported on your system</p> <p>If you specify a maximum file size, you also must specify a maximum number of trace files with the files option.</p> <p>world-readable—(Optional) Allow any user to read the log file.</p> |

| | |
|------------------------------|--|
| Required Privilege | interface—To view this statement in the configuration. |
| Level | interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Tracing APPID Operations on page 686 |

traceoptions (Services Local Policy Decision Function)

| | |
|----------------------------|--|
| Syntax | <pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>>; flag <i>flag</i>; no-remote-trace; } </pre> |
| Hierarchy Level | [edit services local-policy-decision-function],
[edit system services local-policy-decision-function] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Configure local policy decision function (L-PDF) tracing options. |
| Options | <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>, then <i>trace-file.1</i>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p><i>flag</i>—Tracing operation to perform. To specify more than one flag, include multiple flag statements.</p> <ul style="list-style-type: none"> • all—Everything • configuration—Configuration traces • database—Database traces • general—Miscellaneous traces • gres—Graceful Routing Engine switchover (GRES) traces • ptsp-statistics—PTSP statistics traces • rtsock—Routing socket traces • statistics—Statistics traces • subscriber—Subscriber traces <p>no-remote-trace—Disable remote tracing.</p> <p>size <i>size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named <i>trace-file</i> reaches this size, it is renamed <i>trace-file.0</i>. When the <i>trace-file</i> again reaches its maximum size, <i>trace-file.0</i> is renamed</p> |

trace-file.1 and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10240 through 1073741824 or the maximum file size supported on your system

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

| | |
|---------------------------------|---|
| Required Privilege Level | routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Tracing L-PDF Operations on page 707 |

type

| | |
|---------------------------------|---|
| Syntax | <code>type type;</code> |
| Hierarchy Level | [edit services application-identification application <i>application-name</i>]
[edit services application-identification nested-application <i>name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define type of application, such as HTTP or FTP. |
| Options | type —Application type. This is a mandatory value and has a maximum length of 32 characters. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Defining an Application Identification on page 674• Application Identification for Nested Applications on page 681 |

type-of-service

| | |
|---------------------------------|--|
| Syntax | <code>type-of-service <i>service-type</i>;</code> |
| Hierarchy Level | [edit services application-identification application <i>application-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define the type of service by service objective. There is no default value. |
| Options | <p>The following <i>service-type</i> options are available:</p> <ul style="list-style-type: none"> • maximize-reliability—Service designed for maximum reliability in packet transmission. • maximize-throughput—Service designed for maximum throughput. • minimize-delay—Service designed for minimum delay in packet transmission. • minimize-monetary-cost—Service designed for minimum monetary cost. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Defining an Application Identification on page 674 |

url

| | |
|---------------------------------|---|
| Syntax | <code>url <i>url</i>;</code> |
| Hierarchy Level | [edit services application-identification download] |
| Release Information | Statement introduced in Junos OS Release 9.5. |
| Description | Define the URL for application package downloads. |
| Options | <i>url</i> —Download URL. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Automatic Download of Application Package Updates on page 686 |

Link and Multilink Services Configuration Statements

- [acknowledge-retries on page 1597](#)
- [acknowledge-timer on page 1597](#)
- [action-red-differential-delay on page 1598](#)
- [address \(Interfaces\) on page 1598](#)

- [bundle](#) on page 1599
- [destination \(Interfaces\)](#) on page 1600
- [disable-mlppp-inner-ppp-pfc](#) on page 1601
- [dlci](#) on page 1601
- [drop-timeout](#) on page 1602
- [encapsulation \(Logical Interface\)](#) on page 1603
- [encapsulation \(Physical Interface\)](#) on page 1604
- [family](#) on page 1605
- [fragment-threshold](#) on page 1606
- [hello-timer](#) on page 1606
- [interfaces](#) on page 1607
- [interleave-fragments](#) on page 1607
- [lmi-type](#) on page 1608
- [minimum-links](#) on page 1608
- [mlfr-uni-nni-bundle-options](#) on page 1609
- [mrru](#) on page 1610
- [mtu](#) on page 1611
- [multicast-dlci](#) on page 1611
- [n391](#) on page 1612
- [n392](#) on page 1613
- [n393](#) on page 1614
- [red-differential-delay](#) on page 1614
- [short-sequence](#) on page 1615
- [t391](#) on page 1615
- [t392](#) on page 1616
- [unit \(Interfaces\)](#) on page 1617
- [yellow-differential-delay](#) on page 1618

acknowledge-retries

| | |
|---------------------------------|--|
| Syntax | <code>acknowledge-retries <i>number</i>;</code> |
| Hierarchy Level | [edit interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services interfaces only, configure the number of retransmission attempts to be made for consecutive hello or remove link messages following the expiration of the acknowledgment timer. |
| Options | <p><i>number</i>—Number of retransmission attempts to be made following the expiration of the acknowledgment timer.</p> <p>Range: 1 through 5</p> <p>Default: 2</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • action-red-differential-delay on page 1598, hello-timer on page 1606 • Configuring Acknowledgment Timers on Link Services Physical Interfaces on page 727 |

acknowledge-timer

| | |
|---------------------------------|---|
| Syntax | <code>acknowledge-timer <i>milliseconds</i>;</code> |
| Hierarchy Level | [edit interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services interfaces only, configure the maximum time, in milliseconds, to wait for an add link acknowledgment, hello acknowledgment, or remove link acknowledgment message. |
| Options | <p><i>milliseconds</i>—Time to wait for an add link acknowledgment, hello acknowledgment, or remove link acknowledgment message.</p> <p>Range: 1 through 10 milliseconds</p> <p>Default: 4 milliseconds</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • address (Interfaces) on page 1598 • hello-timer on page 1606 • Configuring Acknowledgment Timers on Link Services Physical Interfaces on page 727 |

action-red-differential-delay

| | |
|---------------------------------|--|
| Syntax | <code>action-red-differential-delay (disable-tx remove-link);</code> |
| Hierarchy Level | [edit interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services interfaces only, configure the action to be taken when the differential delay exceeds the red limit. |
| Options | disable-tx —Disable transmission on the bundle link.
remove-link —Remove the bundle link from service.
Default: <code>remove-link</code> |
| Required Privilege Level | <code>interface</code> —To view this statement in the configuration.
<code>interface-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• yellow-differential-delay on page 1618• Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16 on page 727 |

address (Interfaces)

| | |
|---------------------------------|---|
| Syntax | <pre>address <i>address</i> {
 destination <i>address</i>;
}</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the interface address. |
| Options | address —Address of the interface.

The remaining statements are explained separately. |
| Required Privilege Level | <code>interface</code> —To view this statement in the configuration.
<code>interface-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces.• Configuring the Links in a Multilink or Link Services Bundle on page 721• <i>Junos OS Network Interfaces Library for Routing Devices</i> |

bundle

| | |
|---------------------------------|--|
| Syntax | <code>bundle (ml-<i>fpc/pic/port</i> ls-<i>fpc/pic/port</i>);</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mlfr-end-to-end],
[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mlfr-uni-nni] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Associate the multilink interface with the logical interface it is joining. |
| Options | ml-<i>fpc/pic/port</i> —Name of the multilink interface you are linking.

ls-<i>fpc/pic/port</i> —Name of the link services interface you are linking. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Links in a Multilink or Link Services Bundle on page 721 |

destination (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>destination address;</code> |
| Hierarchy Level | <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]</code>
<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],</code>
<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]</code>
<code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>For CoS on ATM interfaces, specify the remote address of the connection.</p> <p>For point-to-point interfaces only, specify the address of the interface at the remote end of the connection.</p> <p>For tunnel and encryption interfaces, specify the remote address of the tunnel.</p> |
| Options | address —Address of the remote side of the connection. |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Linear RED Profiles on ATM Interfaces• Multilink and Link Services Logical Interface Configuration Overview on page 717• Configuring Encryption Interfaces on page 1251• Configuring Traffic Sampling on page 871• Configuring Flow Monitoring on page 818• Configuring Unicast Tunnels on page 1213 |

disable-mlppp-inner-ppp-pfc

| | |
|---------------------------------|--|
| Syntax | <code>disable-mlppp-inner-ppp-pfc;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced in Junos OS Release 8.2. |
| Description | For MLPPP interfaces only, disable compression of the inner PPP header in the MLPPP payload. By default, compression is enabled. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces on page 730 |

dlci

| | |
|---------------------------------|--|
| Syntax | <code>dlci <i>dlci-identifier</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>For Frame Relay and Multilink Frame Relay user-to-network interface (UNI) network-to-network interface (NNI) encapsulation only, and for link services and point-to-point interfaces only, configure the data-link connection identifier (DLCI) for a permanent virtual circuit (PVC) or a switched virtual circuit (SVC).</p> <p>To configure a DLCI for a point-to-multipoint interface, use the multipoint-destination statement to specify the DLCI.</p> |
| Options | <p><i>dlci-identifier</i>—Data-link connection identifier.</p> <p>Range: 16 through 1022</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring DLCIs on Link Services Logical Interfaces on page 777 • Junos OS Network Interfaces Library for Routing Devices |

drop-timeout

| | |
|---------------------------------|---|
| Syntax | <code>drop-timeout <i>milliseconds</i>;</code> |
| Hierarchy Level | [edit interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options],
[edit interfaces (<i>ls-fpc/pic/port</i> <i>ml-fpc/pic/port</i>) unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options],
[edit logical-systems <i>logical-system-name</i> interfaces (<i>ls-fpc/pic/port</i> <i>ml-fpc/pic/port</i>) unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For multilink and link services interfaces only, configure the drop timeout period, in milliseconds. |
| Options | <i>milliseconds</i> —Drop timeout period.
Range: 1 through 2000 milliseconds
Default: 500 ms for bundles greater than or equal to the T1 bandwidth value, and 1500 ms for other bundles. Any CLI-configured value overrides these defaults. Setting a value of 0 reverts to the default. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces on page 730 |

encapsulation (Logical Interface)

| | |
|---------------------------------|--|
| Syntax | <code>encapsulation {atm-mlppp-llc multilink-frame-relay-end-to-end multilink-ppp ... };</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Logical link-layer encapsulation type. |
| Options | <p>atm-mlppp-llc—For ATM 2 interfaces, use Multilink Point-to-Point Protocol (MLPPP) over ATM Adaptation Layer 5 (AAL5) logical link control (LLC) encapsulation, as described in RFC 2364, <i>PPP over AAL5</i>.</p> <p>multilink-frame-relay-end-to-end—Use Multilink Frame Relay (MLFR) FRF.15 encapsulation. This encapsulation is used on multilink “link services interfaces and their constituent T1 or E1 interfaces”, and is supported on LSQ and redundant LSQ interfaces.</p> <p>multilink-ppp—Use MLPPP encapsulation. This encapsulation is used only on multilink and link services interfaces and their constituent T1 or E1 interfaces.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 729 • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

encapsulation (Physical Interface)

| | |
|---------------------------------|--|
| Syntax | encapsulation (multilink-frame-relay-uni-nni ...); |
| Hierarchy Level | [edit interfaces <i>interface-name</i>],
[edit interfaces rlsqnumber: <i>number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Physical link-layer encapsulation type. |
| Default | MLFR UNI NNI encapsulation (on link services interfaces). |
| Options | multilink-frame-relay-uni-nni —Use MLFR UNI NNI encapsulation. This encapsulation is used only on link services interfaces functioning as FRF.16 bundles and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Encapsulation for Link Services Physical Interfaces on page 726• <i>Junos OS Network Interfaces Library for Routing Devices</i> |

family

| | |
|---------------------------------|---|
| Syntax | <pre>family <i>family</i> { <i>address</i> <i>address</i> { <i>destination</i> <i>address</i>; } }</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> <i>unit</i> <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure protocol family information for the logical interface. |
| Options | <p><i>family</i>—Protocol family:</p> <ul style="list-style-type: none"> • ccc—Circuit cross-connect protocol suite • inet—IP version 4 (IPv4) • inet6—IP version 6 (IPv6) • iso—Open Systems Interconnection (OSI) International Organization for Standardization (ISO) protocol suite • mlfr-end-to-end—Multilink Frame Relay FRF.15 • mlfr-uni-nni—Multilink Frame Relay FRF.16 • multilink-ppp—Multilink Point-to-Point Protocol • mpls—MPLS • tcc—Translational cross-connect protocol suite • tnp—Trivial Network Protocol • vpls—Virtual private LAN service <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Link and Multilink Services Interfaces Feature Guide for Routing Devices</i> • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

fragment-threshold

| | |
|---------------------------------|---|
| Syntax | <code>fragment-threshold bytes;</code> |
| Hierarchy Level | [edit interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options],
[edit interfaces (<i>ls-fpc/pic/port</i> <i>ml-fpc/pic/port</i>) unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options],
[edit logical-systems <i>logical-system-name</i> interfaces (<i>ls-fpc/pic/port</i> <i>ml-fpc/pic/port</i>) unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For multilink and link services interfaces only, set the fragmentation threshold, in bytes. |
| Options | bytes —Maximum size, in bytes, for multilink packet fragments. Any nonzero value must be a multiple of 64 bytes.
Range: 128 through 16,320 bytes
Default: 0 bytes (no fragmentation) |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 732• Example: Configuring a Multilink Interface with MLPPP on page 747 |

hello-timer

| | |
|---------------------------------|--|
| Syntax | <code>hello-timer milliseconds;</code> |
| Hierarchy Level | [edit interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services interfaces only, configure the rate at which hello messages are sent. A hello message is transmitted after a period defined in milliseconds has elapsed. |
| Options | milliseconds —The rate at which hello messages are sent.
Range: 1 through 180 milliseconds
Default: 10 milliseconds |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Acknowledgment Timers on Link Services Physical Interfaces on page 727• address (Interfaces) on page 1598, acknowledge-timer on page 1597 |

interfaces

| | |
|---------------------------------|---|
| Syntax | interfaces { ... } |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure interfaces on the router. |
| Default | The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

interleave-fragments

| | |
|---------------------------------|--|
| Syntax | interleave-fragments; |
| Hierarchy Level | [edit interfaces ls-fpc/pic/port:channel unit logical-unit-number],
[edit logical-systems logical-system-name interfaces ls-fpc/pic/port unit logical-unit-number] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>For link services and voice services interfaces only, interleave long packets with high-priority packets.</p> <p>Allows small delay-sensitive packets, such as voice over IP (VoIP) packets, to interleave with long fragmented packets. This minimizes the latency of delay-sensitive packets.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces on page 773 |

lmi-type

| | |
|---------------------------------|--|
| Syntax | <code>lmi-type (ansi itu);</code> |
| Hierarchy Level | <code>[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Set the Frame Relay Local Management Interface (LMI) type. |
| Options | ansi —Use American National Standards Institute (ANSI) T1.167 Annex D LMIs.
itu —Use ITU Q933 Annex A LMIs.
Default: <code>itu</code> |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Keepalives on Link Services Physical Interfaces on page 728 |

minimum-links

| | |
|---------------------------------|---|
| Syntax | <code>minimum-links number;</code> |
| Hierarchy Level | <code>[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options],</code>
<code>[edit interfaces (ls-fpc/pic/port ml-fpc/pic/port) unit logical-unit-number],</code>
<code>[edit logical-systems logical-system-name interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options],</code>
<code>[edit logical-systems logical-system-name interfaces (ls-fpc/pic/port ml-fpc/pic/port) unit logical-unit-number]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For multilink or link services interfaces only, set the minimum number of links that must be up for the bundle to be labeled up. A member link is considered up when the PPP Link Control Protocol (LCP) phase transitions to open state.

The minimum-links value should be identical on both ends of the bundle. |
| Options | number —Number of links.
Range: 1 through 8
Default: 1 |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces on page 733 |

mlfr-uni-nni-bundle-options

| | |
|---------------------------------|---|
| Syntax | <pre>mlfr-uni-nni-bundle-options { acknowledge-retries <i>number</i>; acknowledge-timer <i>milliseconds</i>; action-red-differential-delay (disable-tx remove-link); cisco-interoperability send-lip-remove-link-for-link-reject; drop-timeout <i>milliseconds</i>; fragment-threshold <i>bytes</i>; hello-timer <i>milliseconds</i>; lmi-type (ansi itu c-lmi); minimum-links <i>number</i>; mrru <i>bytes</i>; n391 <i>number</i>; n392 <i>number</i>; n393 <i>number</i>; red-differential-delay <i>milliseconds</i>; t391 <i>number</i>; t392 <i>number</i>; yellow-differential-delay <i>milliseconds</i>; }</pre> |
| Hierarchy Level | [edit interfaces <i>ls-fpc/pic/port</i> : <i>channel</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>Configure link services interface management properties.</p> <p>The statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> Configuring Encapsulation for Link Services Physical Interfaces on page 726 |

mrru

| | |
|---------------------------------|---|
| Syntax | <code>mrru bytes;</code> |
| Hierarchy Level | [edit interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options],
[edit interfaces (<i>ml-fpc/pic/port</i> <i>ls-fpc/pic/port</i>) unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options],
[edit logical-systems <i>logical-system-name</i> interfaces (<i>ml-fpc/pic/port</i> <i>ls-fpc/pic/port</i>) unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For multilink or link services interfaces only, set the maximum received reconstructed unit (MRRU). The MRRU is similar to the maximum transmission unit (MTU), but is specific to multilink interfaces. |
| Options | bytes —MRRU size.
Range: 1500 through 4500 bytes
Default: 1500 bytes |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring MRRU on Multilink and Link Services Logical Interfaces on page 734 |

mtu

| | |
|---------------------------------|---|
| Syntax | <code>mtu bytes;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i>],
[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values. |
| Options | bytes —MTU size.
Range: 0 through 5012 bytes
Default: 1500 bytes (inet , inet6 , and iso families), 1448 bytes (mpls) |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring MRRU on Multilink and Link Services Logical Interfaces on page 734 • Junos OS Network Interfaces Library for Routing Devices |

multicast-dlci

| | |
|---------------------------------|--|
| Syntax | <code>multicast-dlci <i>dlci-identifier</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For point-to-multipoint link services interfaces only, enable multicast support on the interface. You can configure multicast support on the interface if the Frame Relay switch performs multicast replication. |
| Options | <i>dlci-identifier</i> —DLCI identifier, a number from 16 through 1022 that defines the Frame Relay DLCI over which the switch expects to receive multicast packets for replication. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Multicast-Capable DLCIs for MLFR FRF.16 Bundles on page 778 |

n391

| | |
|---------------------------------|---|
| Syntax | n391 <i>number</i> ; |
| Hierarchy Level | [edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services interfaces only, set the Frame Relay full status polling interval. |
| Options | <i>number</i> —Polling interval.
Range: 1 through 255
Default: 6 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Keepalives on Link Services Physical Interfaces on page 728• n392 on page 1613• , n393 on page 1614• t391 on page 1615• t392 on page 1616 |

n392

| | |
|---------------------------------|---|
| Syntax | <code>n392 number;</code> |
| Hierarchy Level | [edit interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services interfaces only, set the Frame Relay error threshold, in number of errors. |
| Options | <p><i>number</i>—Error threshold.</p> <p>Range: 1 through 10</p> <p>Default: 3</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Keepalives on Link Services Physical Interfaces on page 728 • n391 on page 1612 • n393 on page 1614 • t391 on page 1615 • t392 on page 1616 |

n393

| | |
|---------------------------------|---|
| Syntax | <code>n393 number;</code> |
| Hierarchy Level | [edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services interfaces only, set the Frame Relay monitored event count. |
| Options | <i>number</i> —Event count.
Range: 1 through 10
Default: 4 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Keepalives on Link Services Physical Interfaces on page 728• n391 on page 1612• n392 on page 1613• t391 on page 1615• t392 on page 1616 |

red-differential-delay

| | |
|---------------------------------|--|
| Syntax | <code>red-differential-delay milliseconds;</code> |
| Hierarchy Level | [edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services interfaces only, configure the red differential delay among bundle links to give warning when a link has a differential delay that exceeds the configured threshold. |
| Options | <i>milliseconds</i> —Red differential delay threshold.
Range: 1 through 2000 milliseconds
Default: 120 milliseconds |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16 on page 727• action-red-differential-delay on page 1598,• yellow-differential-delay on page 1618 |

short-sequence

| | |
|---------------------------------|---|
| Syntax | <code>short-sequence;</code> |
| Hierarchy Level | [edit interfaces (ls-fpc/pic/port ml-fpc/pic/port) unit logical-unit-number],
[edit logical-systems logical-system-name interfaces (ls-fpc/pic/port ml-fpc/pic/port) unit logical-unit-number] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For multilink interfaces only, set the length of the packet sequence identification number to 12 bits. |
| Default | If not included in the configuration, the length is set to 24 bits. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces on page 735 |

t391

| | |
|---------------------------------|---|
| Syntax | <code>t391 number;</code> |
| Hierarchy Level | [edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services interfaces only, set the Frame Relay link integrity polling interval. |
| Options | number —Link integrity polling interval.
Range: 5 through 30 seconds
Default: 10 seconds |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Keepalives on Link Services Physical Interfaces on page 728 • n391 on page 1612 • n392 on page 1613 • t392 on page 1616 • n393 on page 1614 |

t392

| | |
|---------------------------------|---|
| Syntax | t392 <i>number</i> ; |
| Hierarchy Level | [edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services interfaces only, set the Frame Relay polling verification interval. |
| Options | <i>number</i> —Polling verification interval.
Range: 5 through 30 seconds
Default: 15 seconds |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Keepalives on Link Services Physical Interfaces on page 728• n391 on page 1612• n392 on page 1613• n393 on page 1614• t391 on page 1615 |

unit (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <pre> unit logical-unit-number { disable-mlppp-inner-ppp-pfc; dlci dlci-identifier; drop-timeout milliseconds; encapsulation type; fragment-threshold bytes; interleave-fragments; minimum-links number; mrru bytes; multicast-dlci dlci-identifier; short-sequence; family family { address address { destination address; } bundle (ml-fpc/pic/port ls-fpc/pic/port); } } </pre> |
| Hierarchy Level | [edit interfaces interface-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device. |
| Options | <p>logical-unit-number—Number of the logical unit.</p> <p>Range: 0 through 16,384</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Link and Multilink Services Interfaces Feature Guide for Routing Devices</i> • <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces. |

yellow-differential-delay

| | |
|---------------------------------|--|
| Syntax | <code>yellow-differential-delay <i>milliseconds</i>;</code> |
| Hierarchy Level | <code>[edit interfaces ls-<i>fpc/pic/port:channel</i> <i>mlfr-uni-nni-bundle-options</i>]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For link services interfaces only, configure the yellow differential delay among bundle links to give warning when a link has a differential delay that exceeds the configured threshold. |
| Options | <i>milliseconds</i> —Yellow differential delay threshold.
Range: 1 through 2000 milliseconds
Default: 72 milliseconds |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16 on page 727• action-red-differential-delay on page 1598• red-differential-delay on page 1614 |

Monitoring, Sampling, and Collection Services Configuration Statements

- [\[edit forwarding-options\] Hierarchy Level on page 1624](#)
- [\[edit interfaces\] Hierarchy Level on page 1627](#)
- [\[edit services dynamic-flow-control\] Hierarchy Level on page 1628](#)
- [\[edit services flow-collector\] Hierarchy Level on page 1629](#)
- [\[edit services flow-monitoring\] Hierarchy Level on page 1630](#)
- [\[edit services flow-tap\] Hierarchy Level on page 1631](#)
- [\[edit services rpm\] Hierarchy Level on page 1631](#)
- [accounting on page 1634](#)
- [address \(Interfaces\) on page 1635](#)
- [address \(Services Dynamic Flow Capture\) on page 1635](#)
- [aggregate-export-interval on page 1636](#)
- [aggregation on page 1637](#)
- [allowed-destinations on page 1638](#)
- [analyzer-address on page 1638](#)
- [analyzer-id on page 1639](#)
- [archive-sites on page 1639](#)

- [authentication-mode](#) on page 1640
- [autonomous-system-type](#) on page 1641
- [bgp](#) on page 1642
- [capture-group](#) on page 1643
- [cflowd \(Discard Accounting\)](#) on page 1644
- [client-list](#) on page 1645
- [collector](#) on page 1645
- [content-destination](#) on page 1646
- [control-source](#) on page 1647
- [core-dump](#) on page 1648
- [data-fill](#) on page 1649
- [data-format](#) on page 1649
- [data-size](#) on page 1650
- [destination \(Interfaces\)](#) on page 1651
- [destination-interface](#) on page 1652
- [destination-ipv4-address \(RFC 2544 Benchmarking\)](#) on page 1653
- [destination-mac-address \(RFC2544 Benchmarking\)](#) on page 1653
- [destination-port](#) on page 1654
- [destination-udp-port \(RFC 2544 Benchmarking\)](#) on page 1655
- [destinations](#) on page 1655
- [direction \(RFC2544 Benchmarking\)](#) on page 1656
- [disable \(Forwarding Options\)](#) on page 1657
- [dscp-code-point](#) on page 1658
- [duplicates-dropped-periodicity](#) on page 1659
- [dynamic-flow-capture](#) on page 1660
- [engine-id \(Forwarding Options\)](#) on page 1661
- [engine-type](#) on page 1662
- [export-format](#) on page 1663
- [extension-service](#) on page 1664
- [family \(Monitoring\)](#) on page 1665
- [family \(RFC2544 Benchmarking\)](#) on page 1666
- [family \(Sampling\)](#) on page 1667
- [file \(Sampling\)](#) on page 1668
- [file \(Trace Options\)](#) on page 1669
- [file-specification \(File Format\)](#) on page 1669
- [file-specification \(Interface Mapping\)](#) on page 1670
- [filename](#) on page 1670

- [filename-prefix](#) on page 1671
- [files](#) on page 1671
- [filter](#) on page 1672
- [flow-active-timeout](#) on page 1673
- [flow-collector](#) on page 1674
- [flow-export-destination](#) on page 1675
- [flow-export-rate](#) on page 1675
- [flow-inactive-timeout](#) on page 1676
- [flow-server](#) on page 1677
- [flow-table-size](#) on page 1678
- [flow-tap](#) on page 1679
- [ftp \(Flow Collector Files\)](#) on page 1680
- [ftp \(Transfer Log Files\)](#) on page 1681
- [g-duplicates-dropped-periodicity](#) on page 1682
- [g-max-duplicates](#) on page 1683
- [hard-limit](#) on page 1683
- [hard-limit-target](#) on page 1684
- [hardware-timestamp](#) on page 1684
- [history-size](#) on page 1685
- [host-outbound](#) on page 1685
- [udp-tcp-port-swap \(RFC 2544 Benchmarking\)](#) on page 1686
- [in-service \(RFC2544 Benchmarking\)](#) on page 1686
- [inactivity-timeout \(Services RPM\)](#) on page 1687
- [inline-jflow](#) on page 1687
- [input \(Port Mirroring\)](#) on page 1688
- [input \(Sampling\)](#) on page 1688
- [input-interface-index](#) on page 1689
- [input-packet-rate-threshold](#) on page 1689
- [instance \(Sampling\)](#) on page 1690
- [interface \(Accounting or Sampling\)](#) on page 1691
- [interface \(Services Flow Tap\)](#) on page 1692
- [interface-map](#) on page 1692
- [interfaces \(Services Dynamic Flow Capture\)](#) on page 1693
- [interfaces \(Video Monitoring\)](#) on page 1694
- [ip-swap \(RFC 2544 Benchmarking\)](#) on page 1695
- [ipv4-flow-table-size](#) on page 1695
- [ipv4-template](#) on page 1696

- [ipv6-flow-table-size](#) on page 1696
- [ipv6-template](#) on page 1697
- [label-position](#) on page 1697
- [local-dump](#) on page 1698
- [logical-system](#) on page 1698
- [match](#) on page 1699
- [max-connection-duration](#) on page 1699
- [max-duplicates](#) on page 1700
- [max-packets-per-second](#) on page 1701
- [maximum-age](#) on page 1701
- [maximum-connections](#) on page 1702
- [maximum-connections-per-client](#) on page 1703
- [maximum-packet-length](#) on page 1704
- [maximum-sessions](#) on page 1705
- [maximum-sessions-per-connection](#) on page 1706
- [minimum-priority](#) on page 1706
- [mode \(RFC 2544 Benchmarking\)](#) on page 1707
- [monitoring](#) on page 1708
- [moving-average-size](#) on page 1709
- [mpls-ipv4-template](#) on page 1709
- [mpls-template](#) on page 1710
- [multiservice-options](#) on page 1710
- [name-format](#) on page 1711
- [next-hop \(Forwarding Options\)](#) on page 1712
- [next-hop-group \(Forwarding Options\)](#) on page 1713
- [no-filter-check](#) on page 1713
- [no-remote-trace \(Trace Options\)](#) on page 1714
- [no-syslog](#) on page 1714
- [notification-targets](#) on page 1715
- [observation-domain-id](#) on page 1716
- [one-way-hardware-timestamp](#) on page 1717
- [option-refresh-rate](#) on page 1718
- [options-template-id](#) on page 1719
- [output \(Accounting\)](#) on page 1720
- [output \(Monitoring\)](#) on page 1721
- [output \(Port Mirroring\)](#) on page 1722
- [output \(Sampling\)](#) on page 1723

- [output-interface-index](#) on page 1724
- [passive-monitor-mode](#) on page 1724
- [password \(Flow Collector File Servers\)](#) on page 1725
- [password \(Transfer Log File Servers\)](#) on page 1725
- [peer-as-billing-template](#) on page 1726
- [pic-memory-threshold](#) on page 1726
- [pop-all-labels](#) on page 1727
- [port \(Flow Monitoring\)](#) on page 1728
- [port \(RPM\)](#) on page 1728
- [port \(TWAMP\)](#) on page 1729
- [pre-rewrite-tos](#) on page 1729
- [probe](#) on page 1730
- [probe-count](#) on page 1731
- [probe-interval](#) on page 1732
- [probe-limit](#) on page 1732
- [probe-server](#) on page 1733
- [probe-type](#) on page 1734
- [rate \(Forwarding Options\)](#) on page 1735
- [receive-options-packets](#) on page 1735
- [receive-ttl-exceeded](#) on page 1736
- [reflect-mode \(RFC2544 Benchmarking\)](#) on page 1737
- [required-depth](#) on page 1738
- [retry \(Services Flow Collector\)](#) on page 1739
- [retry-delay](#) on page 1739
- [rfc2544-benchmarking](#) on page 1740
- [routing-instance](#) on page 1741
- [routing-instances](#) on page 1742
- [rpm \(Interfaces\)](#) on page 1742
- [rpm \(Services\)](#) on page 1743
- [run-length](#) on page 1745
- [sample-once](#) on page 1745
- [sampling \(Forwarding Options\)](#) on page 1746
- [sampling \(Interfaces\)](#) on page 1748
- [server](#) on page 1749
- [server-inactivity-timeout](#) on page 1749
- [service-port](#) on page 1750
- [service-type \(RFC2544 Benchmarking\)](#) on page 1750

- [services \(RPM\) on page 1751](#)
- [shared-key on page 1751](#)
- [size on page 1752](#)
- [soft-limit on page 1753](#)
- [soft-limit-clear on page 1753](#)
- [source-address \(Forwarding Options\) on page 1754](#)
- [source-address \(Services\) on page 1755](#)
- [source-addresses on page 1755](#)
- [source-id on page 1756](#)
- [source-ipv4-address \(RFC 2544 Benchmarking\) on page 1756](#)
- [source-mac-address \(RFC2544 Benchmarking\) on page 1757](#)
- [source-udp-port \(RFC 2544 Benchmarking\) on page 1757](#)
- [stamp on page 1758](#)
- [syslog on page 1758](#)
- [target \(Services RPM\) on page 1759](#)
- [tcp on page 1759](#)
- [templates on page 1760](#)
- [test on page 1762](#)
- [tests \(RFC 2544 Benchmarking\) on page 1763](#)
- [test-interface \(RFC 2544 Benchmarking\) on page 1764](#)
- [test-interval on page 1765](#)
- [test-name \(RFC 2544 Benchmarking\) on page 1766](#)
- [thresholds on page 1767](#)
- [traceoptions \(Forwarding Options\) on page 1768](#)
- [traceoptions \(RPM\) on page 1769](#)
- [transfer on page 1770](#)
- [transfer-log-archive on page 1771](#)
- [traps on page 1772](#)
- [ttl on page 1773](#)
- [twamp on page 1774](#)
- [twamp-server on page 1775](#)
- [template \(Forwarding Options\) on page 1775](#)
- [template-id on page 1776](#)
- [template-refresh-rate on page 1777](#)
- [trio-flow-offload on page 1777](#)
- [udp on page 1778](#)
- [unit on page 1779](#)

- [username \(Services\)](#) on page 1780
- [variant](#) on page 1780
- [version](#) on page 1781
- [version9 \(Forwarding Options\)](#) on page 1781
- [video-monitoring](#) on page 1782
- [world-readable](#) on page 1783

[edit forwarding-options] Hierarchy Level

To configure flow monitoring and accounting properties, include the following statements at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
accounting name {
  output {
    aggregate-export-interval seconds;
    cflowd hostname {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
    }
    autonomous-system-type (origin | peer);
    port port-number;
    version format;
  }
  flow-active-timeout seconds;
  flow-inactive-timeout seconds;
  interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
  }
}
monitoring name {
  family family {
    output {
      cflowd hostname port port-number;
      export-format format;
      flow-active-timeout seconds;
      flow-export-destination {
        collector-pic;
      }
      flow-inactive-timeout seconds;
      interface interface-name {
        engine-id number;
        engine-type number;
        input-interface-index number;
      }
    }
  }
}
```

```

        output-interface-index number;
        source-address address;
    }
}
}
next-hop-group group-names {
    interface interface-name {
        next-hop address;
    }
}
port-mirroring {
    input {
        rate rate;
        run-length number;
        maximum-packet-length bytes
    }
    family (inet | inet6) {
        output {
            interface interface-name {
                next-hop address;
            }
            no-filter-check;
        }
    }
    traceoptions {
        file filename {
            files number;
            size bytes;
            (world-readable | no-world-readable);
        }
    }
}
sampling {
    disable;
    sample-once;
    input {
        rate number;
        run-length number;
        max-packets-per-second number;
        maximum-packet-length bytes;
    }
    traceoptions {
        no-remote-trace;
        file filename <files number> <size bytes> <match expression> <world-readable |
            no-world-readable>;
    }
}
family (inet | inet6 | mpls) {
    disable;
    output {
        aggregate-export-interval seconds;
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        extension-service service-name;
        flow-server hostname {
            aggregation {
                autonomous-system;
            }
        }
    }
}

```

```

        destination-prefix;
        protocol-port;
        source-destination-prefix {
            caida-compliant;
        }
        source-prefix;
    }
    autonomous-system-type (origin | peer);
    (local-dump | no-local-dump);
    port port-number;
    source-address address;
    version format;
    version9 {
        template template-name;
    }
}
interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
}
file {
    disable;
    filename filename;
    files number;
    size bytes;
    (stamp | no-stamp);
    (world-readable | no-world-readable);
}
}
instance instance-name {
    disable;
    input {
        rate number;
        run-length number;
        max-packets-per-second number;
        maximum-packet-length bytes;
    }
    family (inet | inet6 | mpls) {
        disable;
        output {
            aggregate-export-interval seconds;
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            extension-service service-name;
            flow-server hostname {
                aggregation {
                    autonomous-system;
                    destination-prefix;
                    protocol-port;
                    source-destination-prefix {
                        caida-compliant;
                    }
                    source-prefix;
                }
            }
        }
    }
}

```

```

autonomous-system-type (origin | peer);
(local-dump | no-local-dump);
port port-number;
source-address address;
version format;
version9 {
    template template-name;
}
}
interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
}
inline-jflow {
    source-address address;
    flow-export-rate rate;
}
}
}
}
}

```



NOTE: For the complete [edit forwarding-options] hierarchy, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*. This section documents only the statements used in flow monitoring and accounting services.

- Related Documentation**
- [\[edit interfaces\] Hierarchy Level on page 1627](#)
 - [\[edit services flow-monitoring\] Hierarchy Level on page 1630](#)

[edit interfaces] Hierarchy Level

To configure flow monitoring and accounting interfaces, include the following statements at the [edit interfaces] hierarchy level:

```

[edit interfaces]
mo-fpc/pic/port {
    unit logical-unit-number {
        family inet {
            accounting {
                destination-class-usage;
                source-class-usage direction;
            }
        }
    }
    address address {
        destination address;
    }
    filter {
        group filter-group-number;
        input filter-name;
        output filter-name;
    }
}

```

```
    }
    receive-options-packets;
    receive-ttl-exceeded;
    sampling direction;
  }
}
multiservice-options {
  (core-dump | no-core-dump);
  (syslog | no-syslog);
  flow-control-options {
    down-on-flow-control;
    dump-on-flow-control;
    reset-on-flow-control;
  }
}
(at-fpc/pic/port | fe-fpc/pic/port | ge-fpc/pic/port) {
  passive-monitor-mode;
}
so-fpc/pic/port {
  unit logical-unit-number {
    passive-monitor-mode;
  }
}
```

- Related Documentation**
- [\[edit forwarding-options\] Hierarchy Level on page 1624](#)
 - [\[edit services flow-monitoring\] Hierarchy Level on page 1630](#)

[\[edit services dynamic-flow-control\] Hierarchy Level](#)

To configure dynamic flow capture, include the **dynamic-flow-capture** statement at the **[edit services]** hierarchy level:

```
[edit services]
dynamic-flow-capture {
  capture-group client-name {
    content-destination identifier {
      address address;
      hard-limit bandwidth;
      hard-limit-target bandwidth;
      soft-limit bandwidth;
      soft-limit-clear bandwidth;
      ttl hops;
    }
    control-source identifier {
      allowed-destinations [ destinations ];
      minimum-priority value;
      no-syslog;
      notification-targets address port port-number;
      service-port port-number;
      shared-key value;
      source-addresses [ addresses ];
    }
  }
  duplicates-dropped-periodicity seconds;
  input-packet-rate-threshold rate;
```

```

    interfaces interface-name;
    max-duplicates number;
    pic-memory-threshold percentage percentage;
  }
  g-duplicates-dropped-periodicity seconds;
  g-max-duplicates number;
  traceoptions{
    file filename <files number> <size size> <world-readable | non-world-readable>;
  }
}

```

Related Documentation • [Configuring Junos Capture Vision on page 849](#)

[edit services flow-collector] Hierarchy Level

To configure flow collection, include the **flow-collector** statement at the **[edit services]** hierarchy level:

```

flow-collector {
  analyzer-address address;
  analyzer-id name;
  destinations {
    ftp:url {
      password "password";
    }
  }
  file-specification {
    variant variant-number {
      data-format format;
      name-format format;
      transfer {
        record-level number;
        timeout seconds;
      }
    }
  }
  interface-map {
    collector interface-name;
    file-specification variant-number;
    interface-name {
      collector interface-name;
      file-specification variant-number;
    }
  }
  retry number;
  retry-delay seconds;
  transfer-log-archive {
    archive-sites {
      ftp:url {
        password "password";
        username username;
      }
    }
  }
  filename-prefix prefix;
  maximum-age minutes;
}

```

```
}  
}
```

**Related
Documentation**

- [Configuring Flow Collection on page 840](#)
- [Sending cflowd Records to Flow Collector Interfaces on page 843](#)
- [Configuring Flow Collection Mode and Interfaces on Services PICs on page 844](#)

[edit services flow-monitoring] Hierarchy Level

```
services {  
  flow-monitoring {  
    version9 {  
      template template-name {  
        flow-active-timeout seconds;  
        flow-inactive-timeout seconds;  
        ipv4-template {  
          nexthop-options {  
            mpls {  
              label-position [ positions ];  
            }  
          }  
        }  
        ipv6-template;  
        mpls-template {  
          label-position [ positions ];  
        }  
        mpls-ipv4-template {  
          label-position [ positions ];  
        }  
        option-refresh-rate {  
          packets packets;  
          seconds seconds;  
        }  
        peer-as-billing-template;  
        template-refresh-rate {  
          packets packets;  
          seconds seconds;  
        }  
        peer-as-billing-template;  
        option-refresh-rate packets;  
        template-refresh-rate packets;  
      }  
    }  
  }  
}
```

**Related
Documentation**

- [Notational Conventions Used in Junos OS Configuration Hierarchies](#)
- [\[edit services\] Hierarchy Level](#)

[edit services flow-tap] Hierarchy Level

To configure flow-tap services, include the **flow-tap** statement at the **[edit services]** hierarchy level. You can also specify whether you want to apply the flow-tap service to IPv4 traffic or IPv6 traffic by including the **family inet | inet6** statement. If the **family** statement is not included in the configuration, the flow-tap service is applied only to the IPv4 traffic.

```
flow-tap {
  interface interface-name;
  family inet | inet6;
}
```

Other statements are configured at the **[edit interfaces]** and **[edit system]** hierarchy levels.

Related Documentation

- [Junos Packet Vision Architecture on page 860](#)
- [Configuring Junos Packet Vision on page 861](#)
- [Configuring FlowTapLite on page 864](#)

[edit services rpm] Hierarchy Level

To configure Real-Time Performance Monitoring (RPM) services, include the **rpm** statement at the **[edit services]** hierarchy level:

```
[edit services]
rpm {
  bgp {
    data-fill data;
    data-size size;
    destination-port port;
    history-size size;
    logical-system logical-system-name [routing-instances routing-instance-name];
    moving-average-size number;
    probe-count count;
    probe-interval seconds;
    probe-type type;
    routing-instances instance-name;
    test-interval interval;
  }
  probe owner {
    test test-name {
      data-fill data;
      data-size size;
      destination-interface interface-name;
      destination-port port;
      dscp-code-point dscp-bits;
      hardware-timestamp;
      history-size size;
      moving-average-size number;
      one-way-hardware-timestamp;
      probe-count count;
```

```
    probe-interval seconds;
    probe-type type;
    routing-instance instance-name;
    source-address address;
    target (url url | address address);
    test-interval interval;
    thresholds thresholds;
    traps traps;
}
}
probe-server {
    tcp {
        destination-interface interface-name;
        port number;
    }
    udp {
        destination-interface interface-name;
        port number;
    }
}
probe-limit limit;
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag;
}
twamp {
    server {
        authentication-mode (authenticated | encrypted | none);
        authentication-key-chain identifier {
            key-id identifier {
                secret password-string;
            }
        }
        client-list list-name {
            [ address address ];
        }
        inactivity-timeout seconds;
        maximum-connections-duration hours;
        maximum-connections count;
        maximum-connections-per-client count;
        maximum-sessions count;
        maximum-sessions-per-connection count;
        port number;
        routing-instance-list {
            instance-name {
                port number;
            }
        }
        server-inactivity-timeout minutes;
    }
}
rfc2544-benchmarking {
    tests {
        test-name test-name {
            test-interface interface-name;
        }
    }
}
```

```

mode reflect;
family (inet | ccc);
destination-ipv4-address address;
destination-udp-port port-number;
source-ipv4-address address;
source-udp-port port-number;
direction (egress | ingress);
}
}
}

```



NOTE: RPM does not require an Adaptive Services (AS) or Multiservices PIC or Multiservices Dense Port Concentrator (DPC) unless you are configuring RPM timestamping as described in [“Configuring RPM Timestamping” on page 964](#).

Related Documentation

- [Configuring BGP Neighbor Discovery Through RPM on page 971](#)
- [Configuring RPM Probes on page 959](#)
- [Configuring RPM Receiver Servers on page 963](#)
- [Limiting the Number of Concurrent RPM Probes on page 964](#)
- [Configuring RPM Timestamping on page 964](#)
- [Configuring TWAMP on page 968](#)
- [Enabling RPM for the Junos OS extension-provider package on page 981](#)
- [Tracing RPM Operations on page 975](#)

accounting

Syntax `accounting name {
 output {
 aggregate-export-interval seconds;
 cflowd hostname {
 aggregation {
 autonomous-system;
 destination-prefix;
 protocol-port;
 source-destination-prefix {
 caida-compliant;
 }
 source-prefix;
 }
 }
 autonomous-system-type (origin | peer);
 port port-number;
 version format;
 }
 flow-active-timeout seconds;
 flow-inactive-timeout seconds;
 interface interface-name {
 engine-id number;
 engine-type number;
 source-address address;
 }
 }
 }
 }`

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the discard accounting instance name and options.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Configuring Discard Accounting on page 883](#)

address (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>address address {
 destination address;
}</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the interface address. |
| Options | <p>address—Address of the interface.</p> <p>The remaining statement is explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Junos OS Network Interfaces Library for Routing Devices</i> for other options not associated with flow monitoring. • Configuring Flow Monitoring on page 818 • Configuring Traffic Sampling on page 871 |

address (Services Dynamic Flow Capture)

| | |
|---------------------------------|--|
| Syntax | <code>address address;</code> |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Configure an IP address for the flow capture destination. |
| Options | address —IP address for the content destination. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Content Destination on page 850 |

aggregate-export-interval

| | |
|---------------------------------|--|
| Syntax | <code>aggregate-export-interval <i>seconds</i>;</code> |
| Hierarchy Level | [edit forwarding-options accounting name output],
[edit forwarding-options sampling instance instance-name family (inet inet6 mpls) output],
[edit forwarding-options sampling family (inet inet6 mpls) output] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the duration, in seconds, of the interval for exporting aggregate accounting information. |
| Options | <i>seconds</i> —Duration. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Discard Accounting on page 883 |

aggregation

| | |
|---------------------------------|--|
| Syntax | <pre> aggregation { autonomous-system; destination-prefix; protocol-port; source-destination-prefix { caida-compliant; } source-prefix; } </pre> |
| Hierarchy Level | <p>[edit forwarding-options accounting output cflowd hostname],</p> <p>[edit forwarding-options sampling instance instance-name family (inet inet6 mpls) output flow-server hostname],</p> <p>[edit forwarding-options sampling family (inet inet6 mpls) output flow-server hostname]</p> |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For cflowd version 8 only, specify the type of data to be aggregated; cflowd records and sends only those flows that match the specified criteria. |
| Options | <p>autonomous-system—Aggregate by autonomous system (AS) number.</p> <p>caida-compliant—Record source and destination mask-length values in compliance with the Version 2.1b1 release of CAIDA's cflowd application. If this statement is not configured, the Junos OS records source and destination mask length values in compliance with the <i>cflowd Configuration Guide</i>, dated August 30, 1999.</p> <p>destination-prefix—Aggregate by destination prefix.</p> <p>protocol-port—Aggregate by protocol and port number.</p> <p>source-destination-prefix—Aggregate by source and destination prefix.</p> <p>source-prefix—Aggregate by source prefix.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Enabling Flow Aggregation on page 898 |

allowed-destinations

| | |
|---------------------------------|---|
| Syntax | <code>allowed-destinations [<i>identifiers</i>];</code> |
| Hierarchy Level | <code>[edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>]</code> |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Identify flow capture destinations that are allowed in messages sent from this control source. |
| Options | <i>identifier</i> —Allowed content destination name. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Control Source on page 851 |

analyzer-address

| | |
|---------------------------------|---|
| Syntax | <code>analyzer-address <i>address</i>;</code> |
| Hierarchy Level | <code>[edit services flow-collector]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure an IP address for the packet analyzer that overrides the default value. |
| Options | <i>address</i> —IP address for packet analyzer. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring a Packet Analyzer on page 841 |





analyzer-id

| | |
|---------------------------------|---|
| Syntax | <code>analyzer-id <i>name</i>;</code> |
| Hierarchy Level | [edit services flow-collector] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure an identifier for the packet analyzer that overrides the default value. |
| Options | <i>name</i> —Identifier for packet analyzer. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring a Packet Analyzer on page 841 |

archive-sites

| | |
|---------------------------------|--|
| Syntax | <pre>archive-sites { ftp:url { password "<i>password</i>"; username <i>username</i>; } }</pre> |
| Hierarchy Level | [edit services flow-collector transfer-log-archive] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the destination for transfer logs. |
| Options | The statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Transfer Logs on page 842 |

authentication-mode

| | |
|---------------------------------|--|
| Syntax | authentication-mode (authenticated control-only-encrypted encrypted none); |
| Hierarchy Level | [edit services rpm twamp server]
[edit services rpm twamp client control-connection <i>control-client-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.5.
Statement at the [edit services rpm twamp client control-connection <i>control-client-name</i>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | Specify the authentication or encryption mode support for the TWAMP test protocol. This statement is required in the configuration; if no authentication or encryption is specified, you should set the value to none . |
| Options | <p>authenticated—Data packets are authenticated.</p> <hr/> <p> NOTE: This mode is supported only on TWAMP servers.</p> <hr/> <p>control-only-encrypted—TWAMP control packets are encrypted. TWAMP data packets are in plain text format.</p> <hr/> <p> NOTE: This mode is supported only on TWAMP servers.</p> <hr/> <p>encrypted—Data packets are encrypted.</p> <hr/> <p> NOTE: This mode is supported only on TWAMP servers.</p> <hr/> <p>none—No authentication or encryption.</p> <hr/> <p> NOTE: This mode is supported on both TWAMP servers and clients.</p> <hr/> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring TWAMP on page 968 • Two-Way Active Measurement Protocol Overview |

autonomous-system-type

| | |
|---------------------------------|---|
| Syntax | <code>autonomous-system-type (origin peer);</code> |
| Hierarchy Level | [edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output flow-server <i>hostname</i>],
[edit forwarding-options sampling family (inet inet6 mpls) output flow-server <i>hostname</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the type of AS numbers that cflowd exports. |
| Default | <code>origin</code> |
| Options | <p>origin—Export origin AS numbers of the packet source address in the Source Autonomous System cflowd field.</p> <p>peer—Export peer AS numbers through which the packet passed in the Source Autonomous System cflowd field.</p> |
| Required Privilege Level | <p><code>interface</code>—To view this statement in the configuration.</p> <p><code>interface-control</code>—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Enabling Flow Aggregation on page 898 |

bgp

Syntax `bgp {
 data-fill data;
 data-size size;
 destination-port port;
 history-size size;
 logical-system logical-system-name <routing-instances routing-instance-name>;
 moving-average-size size;
 probe-count count;
 probe-interval seconds;
 probe-type type;
 routing-instances instance-name;
 test-interval interval;
 }`

Hierarchy Level `[edit services rpm bgp]
 [edit protocols bgp group group-name]
 [edit routing-instances instance-name protocols bgp group group-name]
 [edit logical-system logical-system-name protocols bgp group group-name]
 [edit logical-system logical-system-name routing-instances instance-name protocols bgp
 group group-name]`

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure BGP neighbor discovery through Real-Time Performance Monitoring (RPM).

Options **bgp**—Define properties for configuring BGP neighbor discovery.

The remaining statements are explained separately.



NOTE: On MX Series routers, you can configure all the statements. On M Series and T Series routers, you can configure only the **logical-system** and **routing-instances** statements.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring BGP Neighbor Discovery Through RPM on page 971](#)

capture-group

| | |
|---------------------------------|---|
| Syntax | <pre>capture-group <i>client-name</i> { content-destination <i>identifier</i> { address <i>address</i>; hard-limit <i>bandwidth</i>; hard-limit-target <i>bandwidth</i>; soft-limit <i>bandwidth</i>; soft-limit-clear <i>bandwidth</i>; ttl <i>hops</i>; } control-source <i>identifier</i> { allowed-destinations [<i>destinations</i>]; minimum-priority <i>value</i>; no-syslog; notification-targets <i>address</i> port <i>port-number</i>; service-port <i>port-number</i>; shared-key <i>value</i>; source-addresses [<i>addresses</i>]; } duplicates-dropped-periodicity <i>seconds</i>; input-packet-rate-threshold <i>rate</i>; interfaces <i>interface-name</i>; max-duplicates <i>number</i>; pic-memory-threshold <i>percentage percentage</i>; }</pre> |
| Hierarchy Level | [edit services dynamic-flow-capture] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Define the capture group values. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Capture Group on page 849 |

cflowd (Discard Accounting)

| | |
|---------------------------------|--|
| Syntax | <pre>cflowd <i>hostname</i> {
 aggregation {
 autonomous-system;
 destination-prefix;
 protocol-port;
 source-destination-prefix {
 caida-compliant;
 }
 source-prefix;
 }
 autonomous-system-type (origin peer);
 label-position {
 template <i>template-name</i>;
 }
 (local-dump no-local-dump);
 port <i>port-number</i>;
 source-address <i>address</i>;
 version <i>format</i>;
}</pre> |
| Hierarchy Level | [edit forwarding-options accounting name output], |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility cfdcollect.</p> <p>You can configure up to one version 5 and one version 8 flow format at the [edit forwarding-options accounting name output] hierarchy level.</p> |
| Options | <p>hostname—The IP address or identifier of the host system (the workstation running the cflowd utility).</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Enabling Flow Aggregation on page 898 |

client-list

| | |
|---------------------------------|---|
| Syntax | <code>client-list <i>list-name</i> {
 address <i>address</i>;
}</code> |
| Hierarchy Level | [edit services rpm twamp server] |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | List of allowed control client hosts that can connect to this server. Each entry is a Classless Interdomain Routing (CIDR) address (IP address plus mask) that represents a network of allowed hosts. You can configure more than one list, but you must configure at least one client address to enable TWAMP. Each list can contain up to 64 entries. |
| Options | <i>list-name</i> —Name of client address list.

<i>address</i> —Address and mask for an allowed client. |
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring TWAMP on page 968 |

collector

| | |
|---------------------------------|---|
| Syntax | <code>collector <i>interface-name</i>;</code> |
| Hierarchy Level | [edit services flow-collector interface-map] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the default flow collector interface for interface mapping. |
| Options | <i>interface-name</i> —Default flow collector interface. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Interface Mappings on page 842 |


content-destination

| | |
|---------------------------------|--|
| Syntax | <pre>content-destination <i>identifier</i> {
 address <i>address</i>;
 hard-limit <i>bandwidth</i>;
 hard-limit-target <i>bandwidth</i>;
 soft-limit <i>bandwidth</i>;
 soft-limit-clear <i>bandwidth</i>;
 ttl <i>hops</i>;
}</pre> |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i>] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Identify the destination for captured packets. |
| Options | <p><i>identifier</i>—Name of the destination.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Content Destination on page 850 |

control-source

| | |
|---------------------------------|---|
| Syntax | <pre>control-source <i>identifier</i> { allowed-destinations [<i>destinations</i>]; minimum-priority <i>value</i>; no-syslog; notification-targets <i>address</i> port <i>port-number</i>; service-port <i>port-number</i>; shared-key <i>value</i>; source-addresses [<i>addresses</i>]; }</pre> |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i>] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Identify the source of the dynamic flow capture request. |
| Options | <p><i>identifier</i>—Name of control source.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Control Source on page 851 |

core-dump

| | |
|---------------------------------|---|
| Syntax | (core-dump no-core-dump); |
| Hierarchy Level | [edit interfaces mo- <i>fpc/pic/port</i> multiservice-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>A useful tool for isolating the cause of a problem. Core dumping is enabled by default. The directory /var/tmp contains core files. The Junos OS saves the current core file (0) and the four previous core files, which are numbered from 1 through 4 (from newest to oldest):</p> <div> NOTE: By default, all members of a configured user group (with read-only permissions) can access the core dump files and attach them to cases associated with JTAC.</div> |
| | <ul style="list-style-type: none">• core-dump—Enable the core dumping operation.• no-core-dump—Disable the core dumping operation. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Flow Monitoring on page 818 |

data-fill

| | |
|---------------------------------|---|
| Syntax | <code>data-fill <i>data</i>;</code>
<code>data-fill-with-zeros <i>data</i>;</code> |
| Hierarchy Level | [edit services rpm bgp],
[edit services rpm probe owner test <i>test-name</i>]
[edit services rpm twamp client control-connection <i>control-client-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 9.3 for PTX Series Packet Transport Routers.
Statement at the [edit services rpm twamp client control-connection <i>control-client-name</i>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | Specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes. The data-fill statement is not valid with the http-get or http-metadata-get probe types. For TWAMP client, if this knob is set, then fill the test packet with zeros, if the knob is not set then the data content would be random value as indicated in RFC. |
| Options | data —A hexadecimal value; for example, 0-9 , A-F . |
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring BGP Neighbor Discovery Through RPM on page 971 • Configuring RPM Probes on page 959 • Two-Way Active Measurement Protocol Overview |

data-format

| | |
|---------------------------------|---|
| Syntax | <code>data-format <i>format</i>;</code> |
| Hierarchy Level | [edit services flow-collector file-specification variant <i>variant-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the data format for a specific file format variant. |
| Options | format —Data format. Specify flow-compressed as the data format. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring File Formats on page 841 |

data-size

| | |
|---------------------|--|
| Syntax | <code>data-size size;</code> |
| Hierarchy Level | [edit services rpm bgp],
[edit services rpm probe owner test test-name]
[edit services rpm twamp client control-connection control-client-name] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.
Statement at the [edit services rpm twamp client control-connection control-client-name] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | Specify the size of the data portion of ICMP probes. The data-size statement is not valid with the http-get or http-metadata-get probe type. |
| Options | data —The size can be from 0 through 65400
Default: 0 |



NOTE: If you configure the hardware timestamp feature (see [“Configuring RPM Timestamping” on page 964](#)):

- The **data-size** default value is 32 bytes and 32 is the minimum value for explicit configuration. The UDP timestamp probe type is an exception; it requires a minimum data size of 52 bytes.
 - The **data-size** must be at least 100 bytes smaller than the default MTU of the interface of the RPM client interface.
-

| | |
|--------------------------|---|
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring BGP Neighbor Discovery Through RPM on page 971• Two-Way Active Measurement Protocol Overview |

destination (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>destination <i>address</i>;</code> |
| Hierarchy Level | <pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel] [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel] [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]</pre> |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>For CoS on ATM interfaces, specify the remote address of the connection.</p> <p>For point-to-point interfaces only, specify the address of the interface at the remote end of the connection.</p> <p>For tunnel and encryption interfaces, specify the remote address of the tunnel.</p> |
| Options | <i>address</i> —Address of the remote side of the connection. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Linear RED Profiles on ATM Interfaces • Multilink and Link Services Logical Interface Configuration Overview on page 717 • Configuring Encryption Interfaces on page 1251 • Configuring Traffic Sampling on page 871 • Configuring Flow Monitoring on page 818 • Configuring Unicast Tunnels on page 1213 |

destination-interface

| | |
|--------------------------|--|
| Syntax | <code>destination-interface <i>interface-name</i>;</code> |
| Hierarchy Level | [edit services rpm probe owner test <i>test-name</i>],
[edit services rpm probe-server (tcp udp)]
[edit services rpm twamp client control-connection <i>control-client-name</i>] |
| Release Information | Statement introduced in Junos OS Release 7.5.
Statement at the [edit services rpm twamp client control-connection <i>control-client-name</i>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | <p>On M Series and T Series routers, specify a services (sp-) interface that adds a timestamp to RPM probe messages. This feature is supported only with icmp-ping, icmp-ping-timestamp, udp-ping, and udp-ping-timestamp probe types. You must also configure the rpm statement on the sp- interface and include the unit 0 family inet statement with a /32 address.</p> <p>On M Series, MX Series, and T Series routers, specify a multiservices (ms-) interface that adds a timestamp to RPM probe messages. This feature is supported only with icmp-ping, icmp-ping-timestamp, udp-ping, and udp-ping-timestamp probe types. You must also configure the rpm statement on the ms- interface and include the unit 0 family inet statement with a /32 address.</p> <p>To enable RPM for the extension-provider packages on the adaptive services interface, configure the object-cache-size, policy-db-size, and package statements at the [edit chassis fpc <i>slot-number</i> pic <i>pic-number</i> adaptive-services service-package extension-provider] hierarchy level. For the extension-provider package, <i>package-name</i> in the package <i>package-name</i> statement is jservices-rpm.</p> |
| Options | <i>interface-name</i> —Name of the adaptive services interface. |
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring RPM Timestamping on page 964• Configuring RPM Receiver Servers on page 963• Configuring RPM Timestamping on page 964• hardware-timestamp on page 1684• rpm (Interfaces) on page 1742• Enabling RPM for the Junos OS extension-provider package on page 981 |

destination-ipv4-address (RFC 2544 Benchmarking)

| | |
|---------------------------------|---|
| Syntax | <code>destination-ipv4-address <i>address</i>;</code> |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests <i>test-name</i> <i>test-name</i>] |
| Release Information | Statement introduced in Junos OS Release 12.3X52 for ACX Series routers.
Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge Routers. |
| Description | Specify the destination IPv4 address to be used in generated test frames. You must configure this option if you specify <code>inet</code> as the family. This option is not required if you specify <code>cccas</code> the family. |
| Options | <i>address</i> —Valid IPv4 address.
Default: If you do not configure the destination IPv4 address, the default value of 192.168.1.20 is used. |
| Required Privilege Level | <code>interface</code> —To view this statement in the configuration.
<code>interface-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring an RFC 2544-Based Benchmarking Test on page 989 • RFC2544-Based Benchmarking Tests Overview on page 983 • rfc2544-benchmarking on page 1740 |

destination-mac-address (RFC2544 Benchmarking)

| | |
|---------------------------------|---|
| Syntax | <code>destination-mac-address <i>mac-address</i>;</code> |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests <i>test-name</i> <i>test-name</i>] |
| Release Information | Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.
Statement introduced in Junos OS Release 14.2 for MX104 3D Universal Edge Routers. |
| Description | Specify the destination MAC address used in the generated test frames. This is a mandatory parameter for family <code>bridge</code> . |
| Options | <i>mac-address</i> —MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i> —for example, 0011.2233.4455 or 00:11:22:33:44:55. |
| Required Privilege Level | <code>interface</code> —To view this statement in the configuration.
<code>interface-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • rfc2544-benchmarking on page 1740 • RFC2544-Based Benchmarking Tests Overview on page 983 • Configuring an RFC 2544-Based Benchmarking Test on page 989 |

destination-port

| | |
|---------------------------------|---|
| Syntax | <code>destination-port <i>port</i>;</code> |
| Hierarchy Level | [edit services rpm bgp],
[edit services rpm probe owner test <i>test-name</i>]
[edit services rpm twamp client control-connection <i>control-client-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.
Statement at the [edit services rpm twamp client control-connection <i>control-client-name</i>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | <p>Specify the User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) port to which a probe is sent. This statement is used only for TCP or UDP probe types.</p> <p>The value for the destination-port can be only 7 when you configure along with hardware timestamping. A constraint check prevents you for configuring any other value for the destination port in this case.</p> <p>This constraint does not apply when you are using one-way hardware timestamping along with destination-port and either probe-type udp-ping or probe-type udp-ping-timestamp.</p> |
| Options | port —The port number can be 7 or from 49,160 to 65,535. |
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring BGP Neighbor Discovery Through RPM on page 971• Configuring RPM Probes on page 959 |

destination-udp-port (RFC 2544 Benchmarking)

| | |
|---------------------------------|---|
| Syntax | <code>destination-udp-port <i>port-number</i>;</code> |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests <i>test-name</i> <i>test-name</i>] |
| Release Information | Statement introduced in Junos OS Release 12.3X52 for ACX Series routers.
Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge Routers. |
| Description | Specify the UDP port of the destination to be used in the UDP header for the generated frames. If you do not specify the UDP port, the default value of 4041 is used. |
| Options | <i>port-number</i> —UDP port number for the test frames
Default: 4041 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring an RFC 2544-Based Benchmarking Test on page 989 • RFC2544-Based Benchmarking Tests Overview on page 983 • rfc2544-benchmarking on page 1740 |

destinations

| | |
|---------------------------------|--|
| Syntax | <pre>destinations { ftp:url { password "<i>password</i>"; } }</pre> |
| Hierarchy Level | [edit services flow-collector] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the primary and secondary destination FTP servers. |
| Options | The statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Destination FTP Servers for Flow Records on page 840 |

direction (RFC2544 Benchmarking)

| | |
|---------------------------------|---|
| Syntax | direction (egress ingress); |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests test-name test-name] |
| Release Information | Statement introduced in Junos OS Release 12.3X52 for ACX Series routers.
Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge Routers. |
| Description | Specify the direction of the interface on which the test must be run. This parameter is valid only for a ccc family and a bridge family. RFC2544 tests are supported only in the egress direction or the user-to-network interface (UNI) direction of an E-line or E-LAN service parameters in a bridge domain between two routers for unicast traffic. You cannot compute the NNI direction of Ethernet services between two routers for multicast or broadcast traffic. |
| Options | egress —Run the test in the egress direction of the interface (network-to-network interface (NNI)). This option is applicable for a ccc and bridge family.

ingress —Run the test in the ingress direction of the interface (user-to-network interface (UNI)). You cannot configure this option for a bridge family. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• rfc2544-benchmarking on page 1740• RFC2544-Based Benchmarking Tests Overview on page 983• Configuring an RFC 2544-Based Benchmarking Test on page 989 |

disable (Forwarding Options)

| | |
|---------------------------------|--|
| Syntax | disable; |
| Hierarchy Level | [edit forwarding-options port-mirror],
[edit forwarding-options port-mirror instance <i>instance-name</i>],
[edit forwarding-options sampling],
[edit forwarding-options sampling instance <i>instance-name</i>],
[edit forwarding-options sampling family (inet inet6 mpls)],
[edit forwarding-options sampling family (inet inet6 mpls) output file] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement added to port-mirror hierarchy in Junos OS Release 9.6. |
| Description | Disable traffic accounting, port mirroring, or sampling. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Traffic Sampling on page 871• Configuring Port Mirroring on page 931 |

dscp-code-point

| | |
|----------------------------|--|
| Syntax | <code>dscp-code-point <i>dscp-bits</i>;</code> |
| Hierarchy Level | [edit services rpm probe owner test <i>test-name</i>]
[edit services rpm twamp client control-connection <i>control-client-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release for PTX Series Packet Transport Routers.
Statement at the [edit services rpm twamp client control-connection <i>control-client-name</i>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | Specify the value of the Differentiated Services (DiffServ) field within the IP header. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern. |
| Options | <p><i>dscp-bits</i>—A valid 6-bit pattern; for example, 001111, or one of the following configured DSCP aliases:</p> <ul style="list-style-type: none">• af11—Default: 001010• af12—Default: 001100• af13—Default: 001110• af21—Default: 010010• af22—Default: 010100• af23 —Default: 010110• af31 —Default: 011010• af32 —Default: 011100• af33 —Default: 011110• af41 —Default: 100010• af42 —Default:100100• af43 —Default:100110• be—Default: 000000• cs1—Default: 001000• cs2—Default: 010000• cs3—Default: 011000• cs4—Default: 100000• cs5—Default: 101000• cs6—Default: 110000• cs7—Default: 111000 |

- **ef**—Default: 101110
- **nc1**—Default: 110000
- **nc2**—Default: 111000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring RPM Probes on page 959](#)
- [Two-Way Active Measurement Protocol Overview](#)

duplicates-dropped-periodicity

Syntax `duplicates-dropped-periodicity seconds;`

Hierarchy Level [edit services dynamic-flow-capture **capture-group** *client-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the frequency for sending notifications to affected control sources when transmission of duplicate sets of data is restricted because the **max-duplicates** threshold has been reached.

Options ***seconds***—Period for sending DuplicatesDropped notifications.
Default: 30 seconds

Usage Guidelines See .

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [g-duplicates-dropped-periodicity on page 1682](#)
- [Limiting the Number of Duplicates of a Packet on page 855](#)
- [max-duplicates on page 1700](#)

dynamic-flow-capture

Syntax `dynamic-flow-capture {
 capture-group client-name {
 content-destination identifier {
 address address;
 hard-limit bandwidth;
 hard-limit-target bandwidth;
 soft-limit bandwidth;
 soft-limit-clear bandwidth;
 ttl hops;
 }
 control-source identifier {
 allowed-destinations [destinations];
 minimum-priority value;
 no-syslog;
 notification-targets address port port-number;
 service-port port-number;
 shared-key value;
 source-addresses [addresses];
 }
 duplicates-dropped-periodicity seconds;
 input-packet-rate-threshold rate;
 interfaces interface-name;
 max-duplicates number;
 pic-memory-threshold percentage percentage;
 }
 g-duplicates-dropped-periodicity seconds;
 g-max-duplicates number;
 }
 }`

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 7.4.

Description Define the dynamic flow capture properties to be applied to traffic.

Options The remaining statements are explained separately.


Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • *Junos Capture Vision*

engine-id (Forwarding Options)

| | |
|---------------------------------|---|
| Syntax | <code>engine-id <i>number</i>;</code> |
| Hierarchy Level | <code>[edit forwarding-options accounting name output interface <i>interface-name</i>],</code>
<code>[edit forwarding-options monitoring name output interface <i>interface-name</i>],</code>
<code>[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output</code>
<code>interface <i>interface-name</i>],</code>
<code>[edit forwarding-options sampling family (inet inet6 mpls) output interface <i>interface-name</i>]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the engine ID number for flow monitoring and accounting services. |
| Options | <i>number</i> —Identity of accounting interface. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Traffic Sampling on page 871 • Configuring Flow Monitoring on page 818 • Configuring Discard Accounting on page 883 |


engine-type

| | |
|--|---|
| Syntax | <code>engine-type <i>number</i>;</code> |
| Hierarchy Level | <code>[edit forwarding-options accounting name output interface <i>interface-name</i>],</code>
<code>[edit forwarding-options monitoring name output interface <i>interface-name</i>],</code>
<code>[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output</code>
<code>interface <i>interface-name</i>],</code>
<code>[edit forwarding-options sampling family (inet inet6 mpls) output interface <i>interface-name</i>]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>Specify the engine type number for flow monitoring and accounting services. The engine type attribute refers to the type of the flow switching engine, such as the route processor or a line module. The configured engine type is inserted in output cflowd packets. The Source ID, a 32-bit value to ensure uniqueness for all flows exported from a particular device, is the equivalent of the engine type and the engine ID fields.</p> |
| <div> NOTE: You must configure a source address in the output interface statements. The interface-level statement of engine-type is added automatically but you may override this value with manually configured statements to track different flows with a single cflowd collector.</div> | |
| Options | <i>number</i> —Platform-specific accounting interface type. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Traffic Sampling on page 871• Configuring Flow Monitoring on page 818• Configuring Discard Accounting on page 883 |

export-format

| | |
|---------------------------------|--|
| Syntax | <code>export-format <i>format</i>;</code> |
| Hierarchy Level | [edit forwarding-options monitoring name output] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Flow monitoring export format. |
| Options | <i>format</i> —Format of the flows.
Values: 5 or 8
Default: 5 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• version on page 1781• Exporting Flows on page 821 |

extension-service

| | |
|---------------------------------|--|
| Syntax | <code>extension-service <i>service-name</i> {
 <i>provider-specific rules</i>;
}</code> |
| Hierarchy Level | [edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6) output]
[edit forwarding-options sampling family (inet inet6) output]
[edit services service-set <i>service-set-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.0. |
| Description | <p>Define a customer specific sampling configuration.</p> <p>Define a service set or traffic monitoring for applications using application-specific configuration guidelines.</p> |
| | <div> NOTE: If the <code>extension-service</code> statement is specified while configuring a service set, the <code>service-order</code> statement is mandatory.</div> |
| Options | <p><i>provider-specific rules</i>—Provider-specific subhierarchy for services and service sets. See the application-specific documentation for details.</p> <p><i>service-name</i>—Name of the service.</p> |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• service-order• sampling on page 1746 |

family (Monitoring)

```
Syntax  family inet {
        output {
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            export-format format;
            cflowd hostname {
                aggregation {
                    autonomous-system;
                    destination-prefix;
                    protocol-port;
                    source-destination-prefix {
                        caida-compliant;
                    }
                    source-prefix;
                }
            }
            port port-number;
        }
        interface interface-name {
            engine-id number;
            engine-type number;
            input-interface-index number;
            output-interface-index number;
            source-address address;
        }
    }
```

Hierarchy Level [edit forwarding-options **monitoring** *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify input and output interfaces and properties for flow monitoring. Only IPv4 (**inet**) is supported.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Flow Monitoring on page 818](#)

family (RFC2544 Benchmarking)

| | |
|---------------------------------|--|
| Syntax | family (bridge ccc inet); |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests test-name test-name] |
| Release Information | Statement introduced in Junos OS Release 12.3X52 for ACX Series routers.
Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge Routers.
bridge option introduced in Junos OS Release 12.3X53 for ACX Series routers.
bridge option introduced in Junos OS Release 14.2 for MX104 3D Universal Edge Routers. |
| Description | Configure the address type family for the benchmarking test. |
| Options | inet —Run the test on an IPv4 service.

ccc —Run the test on a circuit cross-connect (CCC) or Ethernet pseudowire service. You can run the RFC2544-based benchmarking test either in the egress or ingress direction.

bridge —Indicates that the test is run on a Layer 2 Ethernet line (E- Line) or an Ethernet LAN (E-LAN) service configured in a bridge domain. You can run the RFC2544-based benchmarking test only in the egress direction or the user-to-network interface (UNI) direction of an Ethernet line. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• rfc2544-benchmarking on page 1740• Configuring an RFC 2544-Based Benchmarking Test on page 989• RFC2544-Based Benchmarking Tests Overview on page 983 |

family (Sampling)

```
Syntax  family (inet | inet6 | mpls) {
        disable;
        output {
            aggregate-export-interval seconds;
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            extension-service service-name;
            flow-server hostname {
                aggregation {
                    autonomous-system;
                    destination-prefix;
                    protocol-port;
                    source-destination-prefix {
                        caida-compliant;
                    }
                    source-prefix;
                }
                autonomous-system-type (origin | peer);
                (local-dump | no-local-dump);
                port port-number;
                source-address address;
                version format;
                version9 {
                    template template-name;
                }
            }
            interface interface-name {
                engine-id number;
                engine-type number;
                source-address address;
            }
            file {
                disable;
                filename filename;
                files number;
                size bytes;
                (stamp | no-stamp);
                (world-readable | no-world-readable);
            }
            inline-jflow {
                source-address address;
                flow-export-rate rate;
            }
        }
    }
```

Hierarchy Level [edit forwarding-options [sampling](#)],
[edit forwarding-options [sampling instance](#) *instance-name*]

Release Information Statement introduced before Junos OS Release 7.4.
mpls option introduced in Release 8.3.
inet6 option introduced in Release 9.4.

Description Configure the protocol family to be sampled. IPv4 (**inet**) is supported for most purposes, but you can configure **family mpls** to collect and export MPLS label information or **family inet6** to collect and export IPv6 traffic using flow aggregation version 9.

The remaining statements are explained separately.



NOTE: The `inline-jflow` statement is valid only under the `[edit forwarding-options sampling instance instance-name family inet output]` hierarchy level. The `file` statement is valid only under the `[edit forwarding-options sampling family inet output]` hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation • [Configuring Traffic Sampling on page 871](#)

file (Sampling)

Syntax `file {
 disable;
 filename filename;
 files number;
 size bytes;
 (stamp | no-stamp);
 (world-readable | no-world-readable);
}`

Hierarchy Level [edit forwarding-options **sampling family inet output**]

Release Information Statement introduced before Junos OS Release 7.4.

Description Collect the traffic samples in a file.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation • [Configuring Traffic Sampling on page 871](#)

file (Trace Options)

| | |
|---------------------------------|--|
| Syntax | file <i>filename</i> <files <i>number</i> <size <i>bytes</i> > <world-readable no-world-readable>; |
| Hierarchy Level | [edit forwarding-options port-mirroring traceoptions],
[edit forwarding-options sampling traceoptions] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure information about the files that contain trace logging information. |
| Options | <i>filename</i> —The name of the file containing the trace information.
Default: /var/log/sampled

The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Tracing Traffic Sampling Operations on page 878 |

file-specification (File Format)

| | |
|---------------------------------|--|
| Syntax | file-specification {
variant <i>variant-number</i> {
data-format <i>format</i> ;
name-format <i>format</i> ;
transfer {
record-level <i>number</i> ;
timeout <i>seconds</i> ;
}
}
} |
| Hierarchy Level | [edit services flow-collector] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the file format for the flow collection files. |
| Options | The statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring File Formats on page 841 |

file-specification (Interface Mapping)

| | |
|---------------------------------|---|
| Syntax | file-specification {
variant <i>variant-number</i> ;
} |
| Hierarchy Level | [edit services flow-collector interface-map] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the default file specification for interface mapping. |
| Options | <i>variant-number</i> —Default file format variant. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

filename

| | |
|---------------------------------|--|
| Syntax | filename <i>filename</i> ; |
| Hierarchy Level | [edit forwarding-options sampling family (inet inet6 mpls) output file] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the name of the output file. |
| Options | <i>filename</i> —Name of the file in which to place the traffic samples. All files are placed in the directory <code>/var/tmp</code> . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Traffic Sampling on page 871 |

filename-prefix

| | |
|---------------------------------|---|
| Syntax | <code>filename-prefix <i>prefix</i>;</code> |
| Hierarchy Level | [edit services flow-collector transfer-log-archive] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the filename prefix for log files. |
| Options | <i>prefix</i> —Filename identifier. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Transfer Logs on page 842 |



files

| | |
|---------------------------------|---|
| Syntax | <code>files <i>number</i>;</code> |
| Hierarchy Level | [edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family (inet inet6 mpls) output file],
[edit forwarding-options sampling traceoptions file] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the total number of files to be saved with samples or trace data. |
| Options | <p><i>number</i>—Maximum number of traffic sampling or trace log files. When a file named <i>sampling-file</i> reaches its maximum size, it is renamed <i>sampling-file.0</i>, then <i>sampling-file.1</i>, and so on, until the maximum number of traffic sampling files is reached. Then the oldest sampling file is overwritten.</p> <p>Range: 1 through 100 files</p> <p>Default: 5 files for sampling output; 10 files for trace log information</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Port Mirroring on page 931 • Configuring Traffic Sampling on page 871 |

filter

| | |
|---------------------------------|--|
| Syntax | <pre>filter {
 input <i>filter-name</i>;
 output <i>filter-name</i>;
 group <i>filter-group-number</i>;
}</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Apply a firewall filter to an interface. You can also use filters for encrypted traffic. |
| Options | <p>group <i>filter-group-number</i>—Define an interface to be part of a filter group. The default filter group number is 0.</p> <p>input <i>filter-name</i>—Name of one filter to evaluate when packets are received on the interface.</p> <p>output <i>filter-name</i>—Name of one filter to evaluate when packets are transmitted on the interface.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> or the <i>Junos OS Administration Library for Routing Devices</i>• Configuring Flow Monitoring on page 818 |

flow-active-timeout

| | |
|---|--|
| Syntax | <code>flow-active-timeout seconds;</code> |
| Hierarchy Level | <p>[edit forwarding-options accounting name output],
 [edit forwarding-options monitoring name output],
 [edit forwarding-options sampling instance instance-name family (inet inet6 mpls) output],
 [edit forwarding-options sampling family (inet inet6 mpls) output],
 [edit services flow-monitoring version9]
 [edit services flow-monitoring version-ipfix template template-name]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.
 Support at the [edit services flow-monitoring version-ipfix template <i>template-name</i>] hierarchy level added in Junos OS Release 10.2.</p> |
| Description | Set the interval after which an active flow is exported. |
| <div>  <p>NOTE: The router must include an Adaptive Services, Multiservices, or Monitoring Services PIC for this statement to take effect.</p> </div> | |
| Options | <p>seconds—Duration of the timeout period.</p> <p>Range: 60 through 1800 seconds (for forwarding-options configurations); 10 through 600 seconds (for services configurations)</p> <p>Default: 1800 seconds (for forwarding-options configurations); 60 seconds (for services configurations)</p> |
| <div>  <p>NOTE: In active flow monitoring, the cflowd or flow monitoring version 9 records are exported after a time period that is a multiple of 60 seconds and greater than or equal to the configured active timeout value. For example, if the active timeout value is 90 seconds, the cflowd or flow monitoring version 9 records are exported at 120-second intervals. If the active timeout value is 150 seconds, the cflowd or flow monitoring version 9 records are exported at 180-second intervals, and so forth.</p> </div> | |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Time Periods when Flow Monitoring is Active and Inactive on page 821 • Configuring the Version 9 Template Properties on page 904 |

flow-collector

Syntax flow-collector {
 analyzer-address *address*;
 analyzer-id *name*;
 destinations {
 ftp:url {
 password "*password*";
 }
 }
 file-specification {
 variant *variant-number* {
 data-format *format*;
 name-format *format*;
 transfer {
 record-level *number*;
 timeout *seconds*;
 }
 }
 }
 interface-map {
 collector *interface-name*;
 file-specification *variant-number*;
 interface-name {
 collector *interface-name*;
 file-specification *variant-number*;
 }
 }
 retry *number*;
 retry-delay *seconds*;
 transfer-log-archive {
 archive-sites {
 ftp:url {
 password "*password*";
 username *username*;
 }
 }
 filename-prefix *prefix*;
 maximum-age *minutes*;
 }
 }
 }

Hierarchy Level [edit services]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the flow collection.

Options The statements are explained separately.

Required Privilege interface—To view this statement in the configuration.
 Level interface-control—To add this statement to the configuration.

Related Documentation • [Flow Collection](#)

flow-export-destination


| | |
|---------------------------------|---|
| Syntax | flow-export-destination {
(cflowd-collector collector-pic);
} |
| Hierarchy Level | [edit forwarding-options monitoring <i>group-name</i> family inet output] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure flow collection. |
| Options | cflowd-collector —cflowd collector.

collector-pic —Collector PIC. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | • Exporting Flows on page 821 |

flow-export-rate

| | |
|---------------------------------|---|
| Syntax | flow-export-rate <i>rate</i> ; |
| Hierarchy Level | [edit forwarding-options sampling instance <i>instance-name</i> family inet output inline-jflow] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the flow export rate of monitored packets in kpps. |
| Options | rate —Flow export rate of monitored packets in kpps (from 1 to 400). |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | • Configuring Discard Accounting on page 883
• Configuring Flow Monitoring on page 818
• Configuring Traffic Sampling on page 871 |

flow-inactive-timeout

| | |
|--|---|
| Syntax | <code>flow-inactive-timeout <i>seconds</i>;</code> |
| Hierarchy Level | <code>[edit forwarding-options accounting name output],</code>
<code>[edit forwarding-options monitoring name output],</code>
<code>[edit forwarding-options sampling instance instance-name family (inet inet6 mpls) output],</code>
<code>[edit forwarding-options sampling family (inet inet6 mpls) output],</code>
<code>[edit services flow-monitoring version9]</code>
<code>[edit services flow-monitoring version-ipfix template template-name]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4.
Support at the <code>[edit services flow-monitoring version-ipfix template <i>template-name</i>]</code> hierarchy level added in Junos OS Release 10.2. |
| Description | Set the interval of inactivity that marks a flow inactive. |
| <div> NOTE: The router must include an Adaptive Services, Multiservices, or Monitoring Services PIC for this statement to take effect.</div> | |
| Options | <i>seconds</i> —Duration of the timeout period.
Range: 60 through 1800 seconds (for forwarding-options configurations); 10 through 600 seconds (for services configurations)
Default: 1800 seconds (for forwarding-options configurations); 60 seconds (for services configurations) |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Time Periods when Flow Monitoring is Active and Inactive on page 821• Configuring the Version 9 Template Properties on page 904 |

flow-server

| | |
|----------------------------|---|
| Syntax | <pre> flow-server <i>hostname</i> { aggregation { autonomous-system; destination-prefix; protocol-port; source-destination-prefix { caida-compliant; } source-prefix; } autonomous-system-type (origin peer); (local-dump no-local-dump); port <i>port-number</i>; source-address <i>address</i>; version <i>format</i>; version9 { template <i>template-name</i>; } } </pre> |
| Hierarchy Level | [edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output],
[edit forwarding-options sampling family (inet inet6 mpls) output] |
| Release Information | Statement introduced before Junos OS Release 7.4.
version9 statement introduced in Junos OS Release 8.3. |
| Description | <p>Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility cfdcollect. Specify a host system to collect sampled flows using the version 9 format.</p> <p>You can configure up to one version 5 and one version 8 flow format at the [edit forwarding-options sampling family (inet inet6 mpls) output flow-server <i>hostname</i>] hierarchy level. For the same configuration, you can specify only either version 9 flow record formats or formats using versions 5 and 8, not both types of formats.</p> |
| Options | <p>hostname—The IP address—IPv4 or IPv6—or identifier of the host system (the workstation either running the cflowd utility or collecting traffic flows using version 9).</p> <p>You can configure only one host system for version 9.</p> |



NOTE: IPv6 configuration for **flow-server** is supported only in Junos OS Release 12.3 and later.

Note that when you configure an IPv6 address for the **flow-server** statement, you must also configure an IPv6 address for the **inline-jflow source-address** statement at the **[edit forwarding-options sampling instance *instance-name* family (inet | inet6 | mpls) output]** hierarchy level.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Traffic Sampling on page 871](#)

flow-table-size

Syntax

```
flow-table-size {  
    ipv4-flow-table-size units;  
    ipv6-flow-table-size units;  
    ipv6-extended-attrib;  
}
```

Hierarchy Level [edit chassis fpc *slot-number* inline-services]

Release Information Statement introduced in Junos OS Release 12.1.
ipv6-extended-attrib option added in Junos OS Release 14.2 for MX Series routers.

Description Configure the size of hash tables for inline services sampling.

Options The remaining statements are defined separately.

flow-tap

| | |
|---------------------------------|--|
| Syntax | <pre>flow-tap { (interface <i>interface-name</i> tunnel-interface <i>interface-name</i> family (inet inet6)); }</pre> |
| Hierarchy Level | [edit services] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | Enable the flow-tap or FlowTapLite application on an interface. FlowTapLite is a lighter version of the flow-tap application that is available on MX Series platforms, M120 routers, and M320 routers with Enhanced III FPCs only. |
| Options | <p>interface <i>interface-name</i>—Specify the interface name for the flow-tap application.</p> <p>tunnel-interface <i>interface-name</i>—Specify the tunnel interface name for the FlowTapLite application.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> [<i>edit services flow-tap</i>] <i>Hierarchy Level</i> Configuring Junos Packet Vision on page 861 |

ftp (Flow Collector Files)

| | |
|----------------------------|---|
| Syntax | <code>ftp:url;</code> |
| Hierarchy Level | [edit services flow-collector destination] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the primary and secondary destination FTP server addresses. |
| Options | <p>url—FTP server address. The URL can include the following macros, typed in braces:</p> <ul style="list-style-type: none">• {%D}—Date• {%T}—Time when the file is created• {%I}—Description string for the logical interface configured using the collector interface-name statement at the [edit services flow-collector interface-map] hierarchy• {%N}—Unique, sequential number for each new file created• {am_pm}—AM or PM• {date}—Current date using the {year} {month} {day} macros• {day}—From 01 through 31• {day_abbrev}—Sun through Sat• {day_full}—Sunday through Saturday• {generation number}—Unique, sequential number for each new file created• {hour_12}—From 01 through 12• {hour_24}—From 00 through 23• {ifalias}—Description string for the logical interface configured using the collector statement at the [edit services flow-collector interface-map] hierarchy• {minute}—From 00 through 59• {month}—From 01 through 12• {month_abbrev}—Jan through Dec• {month_full}—January through December• {num_zone}—From -2359 to +2359; this macro is not supported• {second}—From 00 through 60• {time}—Time the file is created, using the {hour_24} {minute} {second} macros• {time_zone}—Time zone code name of the locale; for example, gmt (this macro is not supported).• {year}—In the format YYYY; for example, 1970 |

- **{year_abbrev}**—From 00 through 99

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation • [Configuring Destination FTP Servers for Flow Records on page 840](#)

ftp (Transfer Log Files)

Syntax `ftp:url;`

Hierarchy Level [edit services flow-collector [transfer-log-archive archive-sites](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the primary and secondary destination FTP server addresses.

Options *url*—FTP server address.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation • [Configuring Transfer Logs on page 842](#)

g-duplicates-dropped-periodicity

| | |
|---------------------------------|---|
| Syntax | g-duplicates-dropped-periodicity <i>seconds</i> ; |
| Hierarchy Level | [edit services dynamic-flow-capture] |
| Release Information | Statement introduced in Junos OS Release 9.2. |
| Description | Specify the frequency for sending notifications to affected control sources when transmission of duplicate sets of data is restricted because the g-max-duplicates threshold has been reached. This setting is applied globally; the duplicates-dropped-periodicity setting applied at the capture-group level overrides the global setting. |
| Default | The default period for sending notifications is 30 seconds. |
| Options | <i>seconds</i> —Period for sending DuplicatesDropped notifications. |
| Usage Guidelines | See . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• duplicates-dropped-periodicity on page 1659• Limiting the Number of Duplicates of a Packet on page 855 |

g-max-duplicates

| | |
|---------------------------------|---|
| Syntax | <code>g-max-duplicates <i>number</i>;</code> |
| Hierarchy Level | [edit services dynamic-flow-capture] |
| Release Information | Statement introduced in Junos OS Release 9.2. |
| Description | Specify the maximum number of content destinations to which DFC PICs can send data from a single input set of packets. Limiting the number of duplicates reduces the load on the PIC. This setting is applied globally; the max-duplicates setting applied at the capture-group level overrides the global setting. |
| Default | If no value is configured, a default setting of 3 is used. |
| Options | <i>number</i> —Maximum number of content destinations.
Range: 1 through 64 |
| Usage Guidelines | See . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • max-duplicates on page 1700 • Limiting the Number of Duplicates of a Packet on page 855 |

hard-limit

| | |
|---------------------------------|---|
| Syntax | <code>hard-limit <i>bandwidth</i>;</code> |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>] |
| Release Information | Statement introduced in Junos OS Release 9.2. |
| Description | Specify a bandwidth threshold at which the dynamic flow capture application begins deleting criteria, until the bandwidth falls below the hard-limit-target value. |
| Options | <i>bandwidth</i> —Hard limit threshold, in bits per second. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • hard-limit-target on page 1684 • Configuring the Content Destination on page 850 |

hard-limit-target

| | |
|--------------------------|---|
| Syntax | <code>hard-limit-target <i>bandwidth</i>;</code> |
| Hierarchy Level | [edit <code>services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i></code>] |
| Release Information | Statement introduced in Junos OS Release 9.2. |
| Description | Specify a bandwidth threshold at which the dynamic flow capture application stops deleting criteria. |
| Options | <i>bandwidth</i> —Target value, in bits per second. |
| Usage Guidelines | See . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• hard-limit on page 1683• Configuring the Content Destination on page 850 |

hardware-timestamp

| | |
|--------------------------|---|
| Syntax | <code>hardware-timestamp;</code> |
| Hierarchy Level | [edit <code>services rpm probe owner test test-name</code>] |
| Release Information | Statement introduced in Junos OS Release 8.1.
Statement applied to MX Series routers in Junos OS Release 10.0.
Statement introduced in Junos OS Release 10.3 for EX Series switches. |
| Description | <p>On MX Series routers, on M-320 routers using the Enhanced Queuing MPC, and on EX Series switches only, enable timestamping of RPM probe messages in the Packet Forwarding Engine host processor. This feature is supported only with <code>icmp-ping</code>, <code>icmp-ping-timestamp</code>, <code>udp-ping</code>, and <code>udp-ping-timestamp</code> probe types.</p> <p>When you configure either <code>probe-type udp-ping</code> or <code>probe-type udp-ping-timestamp</code> along with the <code>hardware-timestamp</code> command, the value for the <code>destination-port</code> can be only 7. A constraint check prevents you for configuring any other value for the destination port in this case.</p> <p>This constraint does not apply when you are configuring <code>one-way-hardware-timestamp</code>.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring RPM Timestamping on page 964 |

history-size

| | |
|---------------------------------|--|
| Syntax | history-size <i>size</i> ; |
| Hierarchy Level | [edit services rpm bgp],
[edit services rpm probe owner test test-name]
[edit services rpm twamp client control-connection <i>control-client-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.
Statement at the [edit services rpm twamp client control-connection <i>control-client-name</i>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | Specify the number of stored history entries. |
| Options | <i>size</i> —A value from 0 to 512.
Default: 50 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring BGP Neighbor Discovery Through RPM on page 971 • Configuring RPM Probes on page 959 • <i>Two-Way Active Measurement Protocol Overview</i> |

host-outbound

| | |
|---------------------------------|---|
| Syntax | host-outbound media-interface; |
| Hierarchy Level | [edit chassis] |
| Release Information | Statement introduced in Junos OS Release 13.2 on MX Series 3D Universal Edge Routers. |
| Description | <p>Enable Layer 2 port mirroring of host-generated outbound packets only on MPCs on MX Series 3D Universal Edge routers.</p> <p>This statement enables all Routing Engine-generated Layer 2 injections to execute egress logical interface filters.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Examples: Layer 2 Port-Mirroring at Multiple Levels of the Chassis</i> • <i>Configuring Port Mirroring</i> • <i>Layer 2 Port Mirroring Overview</i> |

udp-tcp-port-swap (RFC 2544 Benchmarking)

| | |
|---------------------------------|---|
| Syntax | udp-tcp-port-swap; |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests <i>test-name</i> <i>test-name</i>] |
| Release Information | Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.
Statement introduced in Junos OS Release 14.2 for MX104 3D Universal Edge routers. |
| Description | Swaps source and destination UDP ports in the test packets. Only UDP port swap and UDP over IPv4 traffic is supported. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• rfc2544-benchmarking on page 1740• RFC2544-Based Benchmarking Tests Overview on page 983• Configuring an RFC 2544-Based Benchmarking Test on page 989 |

in-service (RFC2544 Benchmarking)

| | |
|---------------------------------|---|
| Syntax | in-service; |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests <i>test-name</i> <i>test-name</i>] |
| Release Information | Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.
Statement introduced in Junos OS Release 14.2 for MX104 3D Universal Edge Routers. |
| Description | <p>Runs the test in the in-service mode. In this mode, while the test is running, the rest of the data traffic sent to and from the UNI port under test on the service are not interrupted. Control protocol packets and control protocol peering are not interrupted.</p> <p>If this mode is not configured, the test runs in the default out-of-service mode. In the out-of-service mode, while the test is running, all the data traffic sent to and from the UNI port under test on the service is interrupted. Control protocol peering is not interrupted whereas control protocol packets such as CFM sessions are interrupted.</p> |
| Default | The default service mode for the reflecting egress interface for an E-LAN service is out-of-service mode. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• rfc2544-benchmarking on page 1740• RFC2544-Based Benchmarking Tests Overview on page 983• Configuring an RFC 2544-Based Benchmarking Test on page 989 |

inactivity-timeout (Services RPM)

| | |
|---------------------------------|--|
| Syntax | <code>inactivity-timeout <i>seconds</i>;</code> |
| Hierarchy Level | <code>[edit services rpm twamp <i>server</i>]</code> |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | Inactivity timeout period, in seconds. |
| Options | <i>seconds</i> —Length of time the session is inactive before it times out.
Default: 1800 seconds |
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring TWAMP on page 968 |

inline-jflow

| | |
|---------------------------------|--|
| Syntax | <pre>inline-jflow { <i>source-address</i> <i>address</i>; <i>flow-export-rate</i> <i>rate</i>; }</pre> |
| Hierarchy Level | <code>[edit forwarding-options <i>sampling instance</i> <i>instance-name</i> <i>family</i> inet <i>output</i>]</code> |
| Release Information | Statement introduced in Junos OS Release 10.2.
Statement introduced in Junos OS Release 14.2 for T4000 routers with Type 5 FPC. |
| Description | Specify inline flow monitoring for traffic from the designated address. |
| Options | <i>address</i> —Source IP address. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Inline Active flow Monitoring on page 890 |

input (Port Mirroring)

| | |
|---------------------------------|--|
| Syntax | <pre>input {
 maximum-packet-length bytes
 rate number;
 run-length number;
}</pre> |
| Hierarchy Level | [edit forwarding-options port-mirroring],
[edit forwarding-options port-mirroring instance <i>instance-name</i>]
[edit forwarding-options port-mirroring family (inet inet6)] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure port mirroring on a logical interface.

The statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Port Mirroring on page 931 |

input (Sampling)

| | |
|---------------------------------|--|
| Syntax | <pre>input {
 max-packets-per-second number;
 rate number;
 run-length number;
 maximum-packet-length bytes;
}</pre> |
| Hierarchy Level | [edit forwarding-options sampling],
[edit forwarding-options sampling instance <i>instance-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure traffic sampling on a logical interface.

The statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Traffic Sampling on page 871 |

input-interface-index

| | |
|---------------------------------|---|
| Syntax | <code>input-interface-index <i>number</i>;</code> |
| Hierarchy Level | [edit forwarding-options monitoring name output interface <i>interface-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify a value for the input interface index that overrides the default supplied by SNMP. |
| Options | <i>number</i> —Input interface index value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Flow Monitoring on page 818 |

input-packet-rate-threshold

| | |
|---------------------------------|---|
| Syntax | <code>input-packet-rate-threshold <i>rate</i>;</code> |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i>] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Specify a packet rate threshold value that triggers a system log warning message. |
| Options | <i>rate</i> —Threshold value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Thresholds on page 854 |

instance (Sampling)

```
Syntax  instance instance-name {
        disable;
        input {
            rate number;
            run-length number;
            max-packets-per-second number;
            maximum-packet-length bytes;
        }
        family (inet | inet6 | mpls) {
            disable;
            output {
                aggregate-export-interval seconds;
                flow-active-timeout seconds;
                flow-inactive-timeout seconds;
                extension-service service-name;
                flow-server hostname {
                    aggregation {
                        autonomous-system;
                        destination-prefix;
                        protocol-port;
                        source-destination-prefix {
                            caida-compliant;
                        }
                        source-prefix;
                    }
                    autonomous-system-type (origin | peer);
                    (local-dump | no-local-dump);
                    port port-number;
                    source-address address;
                    version format;
                    version9 {
                        template template-name;
                    }
                }
            }
            interface interface-name {
                engine-id number;
                engine-type number;
                source-address address;
            }
            inline-jflow {
                source-address address;
                flow-export-rate rate;
            }
        }
    }
```

Hierarchy Level [edit forwarding-options [sampling](#)]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure a sampling instance.

The remaining statements are explained separately.

| | |
|---------------------------------|---|
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Sampling Instance Configuration on page 881 |

interface (Accounting or Sampling)

| | |
|---------------------------------|---|
| Syntax | <pre>interface <i>interface-name</i> { engine-id <i>number</i>; engine-type <i>number</i>; source-address <i>address</i>; }</pre> |
| Hierarchy Level | [edit forwarding-options accounting <i>name</i> output],
[edit forwarding-options sampling family (inet inet6 mpls) output],
[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the output interface for monitored traffic. |
| Options | <i>interface-name</i> —Name of the interface.

The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Discard Accounting on page 883 • Configuring Traffic Sampling on page 871 |

interface (Services Flow Tap)

| | |
|---------------------------------|---|
| Syntax | <code>interface sp-fpc/pic/port.logical-unit-number;</code> |
| Hierarchy Level | [edit services flow-tap] |
| Release Information | Statement introduced in Junos OS Release 8.1. |
| Description | Specify the AS PIC interface used with the flow-tap application. Any AS PIC available in the router can be assigned, and any logical interface on the AS PIC can be used. |
| Options | <i>interface-name</i> —Name of the DFC interface.

You cannot configure flow-tap services on channelized interfaces. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Junos Packet Vision Interface on page 861 |

interface-map

| | |
|---------------------------------|--|
| Syntax | <pre>interface-map {
 collector <i>interface-name</i>;
 file-specification <i>variant-number</i>;
 <i>interface-name</i> {
 collector <i>interface-name</i>;
 file-specification <i>variant-number</i>;
 }
}</pre> |
| Hierarchy Level | [edit services flow-collector] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Match an input interface with a flow collector interface and apply the preset file specifications to the input interface. |
| Options | The statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Interface Mappings on page 842 |

interfaces (Services Dynamic Flow Capture)

| | |
|---------------------------------|---|
| Syntax | <code>interfaces <i>interface-name</i>;</code> |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i>] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Specify the DFC interface used with the control source configured in the same capture group. |
| Options | <i>interface-name</i> —Name of the DFC interface. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the DFC PIC Interface on page 852 |

interfaces (Video Monitoring)

Syntax

```
interfaces {  
  interface-name {  
    family {  
      inet {  
        input-flows {  
          input-flow-name {  
            source-address [ address ];  
            destination-address [ address ];  
            source-port [ port ];  
            destination-port [ port ];  
            template template-name;  
          }  
        }  
        output-flows {  
          output-flow-name {  
            source-address [ address ];  
            destination-address [ address ];  
            source-port [ port ];  
            destination-port [ port ];  
            template template-name;  
          }  
        }  
      }  
    }  
  }  
}
```

Hierarchy Level [edit services [video-monitoring](#)]

Release Information Statement introduced in Junos OS Release 14.1.

Description Define video monitoring for specified input or output flows on selected interfaces.

Options *interface-name*—Name of the interace to monitor.

address—Source or destination IPv4 address or prefix value.

port—Port number.

Range: 0 through 65,535

template-name—Name of the template used to monitor flows on an interface.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.


Related Documentation

- [Configuring Inline Video Monitoring on page 1045](#)

ip-swap (RFC 2544 Benchmarking)

| | |
|---------------------------------|---|
| Syntax | <code>ip-swap;</code> |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests <i>test-name</i> <i>test-name</i>] |
| Release Information | Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.
Statement introduced in Junos OS Release 14.2 for MX Series routers. |
| Description | Swaps source and destination IPv4 addresses. This statement is applicable only for family bridge . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • RFC2544-Based Benchmarking Tests Overview on page 983 • Configuring an RFC 2544-Based Benchmarking Test on page 989 • rfc2544-benchmarking on page 1740 |


ipv4-flow-table-size

| | |
|--|--|
| Syntax | <code>ipv4-flow-table-size <i>units</i>;</code> |
| Hierarchy Level | [edit chassis fpc <i>slot-number</i> inline-services flow-table-size] |
| Description | Configure the size of the IPv4 flow table in units of 256K entries. |
| <div style="display: flex; align-items: center;">  <div> <p>NOTE: Any changes in the configured size of the flow has table sizes initiates an automatic reboot of the FPC.</p> </div> </div> | |
| Options | <p>units—Number of 256K flow entries available for the IPv4 flow table.</p> <p>Range: 1 through 15</p> <p>Default: 15 (3840K)</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Inline Active flow Monitoring on page 890 |

ipv4-template

| | |
|---------------------------------|--|
| Syntax | ipv4-template; |
| Hierarchy Level | [edit services flow-monitoring version9 template <i>template-name</i>]
[edit services flow-monitoringversion-ipfix template <i>template-name</i>] |
| Release Information | Statement introduced in Junos OS Release 8.3.
Support at the [edit services flow-monitoring version-ipfix template <i>template-name</i>] hierarchy level added in Junos OS Release 10.2. |
| Description | Specify that the flow aggregation version 9 template is used only for IPv4 records. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902 |

ipv6-flow-table-size

| | |
|--|--|
| Syntax | ipv6-flow-table-size <i>units</i> ; |
| Hierarchy Level | [edit chassis fpc <i>slot-number</i> inline-services ipv6 flow-table-size] |
| Description | Configure the size of the IPv6 flow table in units of 256K entries. |
| <div> NOTE: Any changes in the configured size of the flow has table sizes initiates an automatic reboot of the FPC.</div> | |
| Options | units —Number of 256K flow entries available for the IPv6 flow table.
Range: 1 through 15
Default: 1K |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Inline Active flow Monitoring on page 890 |

ipv6-template

| | |
|---------------------------------|--|
| Syntax | ipv6-template; |
| Hierarchy Level | [edit services flow-monitoring version9 template template-name]
[edit services flow-monitoringversion-ipfix template template-name] |
| Release Information | Statement introduced in Junos OS Release 9.4.
Support at the [edit services flow-monitoring version-ipfix template <i>template-name</i>] hierarchy level added in Junos OS Release 10.2. |
| Description | Specify that the flow aggregation version 9 template is used only for IPv6 records. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902 |

label-position

| | |
|---------------------------------|--|
| Syntax | label-position [<i>positions</i>]; |
| Hierarchy Level | [edit services flow-monitoring version9 template template-name mpls-ipv4-template],
[edit services flow-monitoring version9 template template-name mpls-template] |
| Release Information | Statement introduced in Junos OS Release 8.3. |
| Description | Specify positions for up to three labels in the active flow monitoring version 9 template. |
| Default | [1 2 3] |
| Options | <i>positions</i> —Numbered positions for the labels. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902 |

local-dump

| | |
|---------------------------------|--|
| Syntax | (local-dump no-local-dump); |
| Hierarchy Level | [edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output flow-server <i>hostname</i>],
[edit forwarding-options sampling family (inet inet6 mpls) output flow-server <i>hostname</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Enable collection of cflowd records in a log file. |
| Options | no-local-dump —Do not dump cflowd records to a log file before exporting.

local-dump —Dump cflowd records to a log file before exporting. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Enabling Flow Aggregation on page 898 |

logical-system

| | |
|---------------------------------|--|
| Syntax | logical-system <i>logical-system-name</i> {
[routing-instances <i>instance-name</i>];
} |
| Hierarchy Level | [edit services rpm bgp] |
| Release Information | Statement introduced in Junos OS Release 7.6. |
| Description | Specify the logical system used by the probes.

The remaining statements are explained separately. |
| Options | logical-system-name —Logical system name. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring BGP Neighbor Discovery Through RPM on page 971 |

match

| | |
|---------------------------------|--|
| Syntax | <code>match <i>expression</i>;</code> |
| Hierarchy Level | [edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling traceoptions file] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Regular expression for lines to be logged for tracing. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Port Mirroring on page 931 • Configuring Traffic Sampling on page 871 |


max-connection-duration

| | |
|---------------------------------|---|
| Syntax | <code>max-connection-duration <i>hours</i>;</code> |
| Hierarchy Level | [edit services rpm twamp server] |
| Release Information | Statement introduced in Junos OS Release 11.1. |
| Description | Specify the maximum time a connection can exist between a client and the server. |
| Options | <i>hours</i> —Number of hours a connection can exist between a client and the server. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring TWAMP on page 968 |

max-duplicates

| | |
|---------------------------------|---|
| Syntax | <code>max-duplicates <i>number</i>;</code> |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.2. |
| Description | Specify the maximum number of content destinations to which the DFC PIC can send data from a single input set of packets. Limiting the number of duplicates reduces the load on the PIC. This setting overrides the globally applied g-max-duplicates setting. |
| Default | If no value is configured, a default setting of 3 is used. |
| Options | <i>number</i> —Maximum number of content destinations.
Range: 1 through 64 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• g-max-duplicates on page 1683• Limiting the Number of Duplicates of a Packet on page 855 |

max-packets-per-second

| | |
|--|--|
| Syntax | <code>max-packets-per-second <i>number</i>;</code> |
| Hierarchy Level | [edit forwarding-options sampling input],
[edit forwarding-options sampling instance instance-name input] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the traffic threshold that must be exceeded before packets are dropped. A value of 0 instructs the Packet Forwarding Engine not to sample any traffic. |
| <div>  NOTE: When you configure active monitoring and specify a Monitoring Services, Adaptive Services, or Multiservices PIC in the output statement, the <code>max-packets-per-second</code> value is ignored. </div> | |
| Options | <i>number</i> —Maximum number of packets per second.
Range: 0 through 65,535
Default: 1000 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> Configuring Traffic Sampling on page 871 |

maximum-age

| | |
|---------------------------------|---|
| Syntax | <code>maximum-age <i>minutes</i>;</code> |
| Hierarchy Level | [edit services flow-collector transfer-log-archive] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Maximum age of transfer log file. |
| Options | <i>maximum-age minutes</i> —Transfer log file age.
Range: 1 through 360 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> Configuring Transfer Logs on page 842 |

maximum-connections

| | |
|----------------------------|--|
| Syntax | <code>maximum-connections <i>count</i>;</code> |
| Hierarchy Level | [edit services rpm twamp server] |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | Configure the maximum number of allowed connections between the server and all control client hosts. |



NOTE: The maximum number of connections between the server and all control client hosts must be greater than or equal to the number of connections between the server and a single controlled client host.

| | |
|---------------------------------|--|
| Options | <code><i>count</i></code> —Maximum number of connections.
Range: 1 through 1000
Default: 64 |
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring TWAMP on page 968 |

maximum-connections-per-client

| | |
|----------------------------|--|
| Syntax | <code>maximum-connections-per-client</code> <i>count</i> ; |
| Hierarchy Level | [edit services rpm twamp server] |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | Configure the maximum number of allowed connections between the server and a single control client host. |



NOTE: The maximum number of connections between the server and all control client hosts must be greater than or equal to the number of connections between the server and a single controlled client host.

| | |
|---------------------------------|--|
| Options | <i>count</i> —Maximum number of connections.
Range: 1 through 500
Default: 64 |
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring TWAMP on page 968 |

maximum-packet-length

| | |
|----------------------------|--|
| Syntax | <code>maximum-packet-length bytes;</code> |
| Hierarchy Level | [edit forwarding-options analyzer analyzer-name input],
[edit forwarding-options port-mirroring input],
[edit forwarding-options port-mirroring instance <i>instance-name</i> input],
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance <i>instance-name</i> input] |
| Release Information | Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
The [edit forwarding-options analyzer analyzer-name input] hierarchy level for MX Series routers introduced in Junos OS Release 14.1. |
| Description | Set the maximum length of the packet used for port mirroring or traffic sampling. Packets with lengths greater than the specified maximum are truncated. |



NOTE: The `maximum-packet-length` statement is not supported on MX80 routers.



NOTE: For MX-Series devices with Modular Port Interface Concentrators (MPCs), when `maximum-packet-length` (clip length) is configured for port-mirrored packets and the mirror-destination interface is a next-hop-group, the clip length would be effective only for the first member interface of the next-hop-group. The mirrored packet copy sent to the rest of the interfaces would not be clipped.


Native analyzer sessions (that is, the [edit forwarding-options analyzer analyzer-name input] hierarchy level for MX Series routers) can be configured without specifying input parameters, which would mean that the instance uses default input values: `rate = 1` and `maximum-packet-length = 0`.

| | |
|----------------|--|
| Options | <i>bytes</i> —Maximum length (in bytes) of the mirrored packet or the sampled packet.
Range: 0 through 9216
Default: 0 |
|----------------|--|

For MX Series routers with Modular Port Concentrators (MPCs), port-mirrored or sampled packets can be truncated (or clipped) to any length in the range of 1 to 255 bytes. Only 1 to 255 are valid values for packet truncation on these devices. For other devices, the range is from 0 to 9216. A `maximum-packet-length` value of zero represents that truncation is disabled, and the entire packet is mirrored or sampled.

| | |
|---------------------------------|--|
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Port Mirroring • Configuring Traffic Sampling |

maximum-sessions

| | |
|--|--|
| Syntax | <code>maximum-sessions count;</code> |
| Hierarchy Level | [edit services rpm twamp server] |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | Configure the maximum number of allowed test sessions the server can have running at one time. |
| <div style="display: flex; align-items: center;">  <div> <p>NOTE: The maximum number of test sessions running on the server at one time must be greater than or equal to the maximum number of sessions the server can open on a single client connection.</p> </div> </div> | |
| Options | <p>count—Maximum number of sessions.</p> <p>Range: 1 through 2048</p> <p>Default: 64</p> |
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring TWAMP on page 968 |

maximum-sessions-per-connection

| | |
|----------------------------|---|
| Syntax | <code>maximum-sessions-per-connection <i>count</i>;</code> |
| Hierarchy Level | [edit services rpm twamp server] |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | Configure the maximum number of allowed sessions the server can open on a single client connection. |



NOTE: The maximum number of test sessions running on the server at one time must be greater than or equal to the maximum number of sessions the server can open on a single client connection.

| | |
|---------------------------------|--|
| Options | <i>count</i> —Maximum number of sessions.
Default: 64 |
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring TWAMP on page 968 |

minimum-priority

| | |
|---------------------------------|---|
| Syntax | <code>minimum-priority <i>value</i>;</code> |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>] |
| Release Information | Statement introduced in Junos OS Release 9.2. |
| Description | Specify the minimum priority for the control source. |
| Options | <i>value</i> —Minimum priority value; if not specified, defaults to 0.
Range: 0 through 254 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Control Source on page 851 |

mode (RFC 2544 Benchmarking)

| | |
|---------------------------------|---|
| Syntax | mode reflect; |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests test-name <i>test-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge Routers. |
| Description | Specify the test mode for the packets that are sent during the benchmarking test. |
| Options | reflect —Causes the test frames to be reflected on the chosen service (IPv4 or Ethernet). |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring an RFC 2544-Based Benchmarking Test on page 989• RFC2544-Based Benchmarking Tests Overview on page 983• rfc2544-benchmarking on page 1740 |

monitoring

Syntax `monitoring name {
 family inet {
 output {
 cflowd hostname port-number;
 export-format cflowd-version-5;
 flow-active-timeout seconds;
 flow-export-destination {
 (cflowd-collector | collector-pic);
 }
 flow-inactive-timeout seconds;
 interface interface-name {
 number;
 engine-type number;
 input-interface-index number;
 output-interface-index number;
 source-address address;
 }
 }
 }
 }`

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the flow monitoring instance name and properties.

 The statements are explained separately.

Required Privilege interface—To view this statement in the configuration.
 Level interface-control—To add this statement to the configuration.

Related • [Configuring Flow Monitoring on page 818](#)
Documentation

moving-average-size

| | |
|---------------------------------|--|
| Syntax | <code>moving-average-size <i>number</i>;</code> |
| Hierarchy Level | [edit services rpm bgp],
[edit services rpm probe owner test <i>test-name</i>]
[edit services rpm twamp client control-connection <i>control-client-name</i>] |
| Release Information | Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.
Statement at the [edit services rpm twamp client control-connection <i>control-client-name</i>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | Enable statistical calculation operations to be performed across a configurable number of the most recent samples. |
| Options | <i>number</i> —Number of samples to be used in calculations.
Range: 0 through 255 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring RPM Probes on page 959 • Two-Way Active Measurement Protocol Overview |

mpls-ipv4-template

| | |
|---------------------------------|--|
| Syntax | <code>mpls-ipv4-template {
 label-position [<i>positions</i>];
}</code> |
| Hierarchy Level | [edit services flow-monitoring version9 template <i>template-name</i>] |
| Release Information | Statement introduced in Junos OS Release 8.3. |
| Description | Specify the flow aggregation version 9 properties for templates that combine IPv4 and MPLS records. The remaining statement is explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902 |

mpls-template

| | |
|---------------------------------|--|
| Syntax | <pre>mpls-template {
 label-position [<i>positions</i>];
}</pre> |
| Hierarchy Level | [edit services flow-monitoring version9 template <i>template-name</i>] |
| Release Information | Statement introduced in Junos OS Release 8.3. |
| Description | Specify the flow aggregation version 9 properties for templates used only for MPLS records. The remaining statement is explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902 |

multiservice-options

| | |
|---------------------------------|---|
| Syntax | <pre>multiservice-options {
 (core-dump no-core-dump);
 (syslog no-syslog);
 flow-control-options {
 down-on-flow-control;
 dump-on-flow-control;
 reset-on-flow-control;
 }
}</pre> |
| Hierarchy Level | [edit interfaces <i>mo-fpc/pic/port</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For flow-monitoring interfaces only, configure multiservice-specific interface properties.

The statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Flow Monitoring on page 818 |

name-format

| | |
|----------------------------|---|
| Syntax | <code>name-format "format";</code> |
| Hierarchy Level | [edit services flow-collector file-specification variant <i>variant-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the name format for a specific file format. The files may include supported macros. Use macros to organize files on the external machine to which they are exported from the collector PIC. |
| Options | <p>format—Specify the filename format, within quotation marks. The name format can include the following macros, typed in braces:</p> <ul style="list-style-type: none"> • {%D}—Date • {%T}—Time when the file is created • {%I}—Description string for the logical interface configured using the collector statement at the [edit services flow-collector interface-map] hierarchy level • {%N}—Unique, sequential number for each new file created • {am_pm}—AM or PM • {date}—Current date using the {year} {month} {day} macros • {day}—From 01 through 31 • {day_abbrev}—Sun through Sat • {day_full}—Sunday through Saturday • {generation number}—Unique, sequential number for each new file created • {hour_12}—From 01 through 12 • {hour_24}—From 00 through 23 • {ifalias}—Description string for the logical interface configured using the collector statement at the [edit services flow-collector interface-map] hierarchy level • {minute}—From 00 through 59 • {month}—From 01 through 12 • {month_abbrev}—Jan through Dec • {month_full}—January through December • {num_zone}—From -2359 through +2359; this macro is not supported • {second}—From 00 through 60 • {time}—Time the file is created, using the {hour_24} {minute} {second} macros • {time_zone}—Time zone code name of the locale; for example, gmt (this macro is not supported). |

- **{year}**—In the format YYYY; for example, 1970
- **{year_abbrev}**—From 00 through 99

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring File Formats on page 841](#)

next-hop (Forwarding Options)

Syntax next-hop *address*;

Hierarchy Level [edit forwarding-options port-mirroring **family** (inet | inet6) **output interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the next-hop address for sending copies of packets to an analyzer.

Options *address*—IP address of the next-hop router.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Port Mirroring on page 931](#)

next-hop-group (Forwarding Options)

| | |
|---------------------------------|--|
| Syntax | <pre>next-hop-group <i>group-name</i> { interface <i>interface-name</i> { next-hop <i>address</i>; } }</pre> |
| Hierarchy Level | [edit forwarding-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>Specify the next-hop address for sending copies of packets to an analyzer.</p> <p>It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.</p> |
| Options | <p><i>address</i>—IP address of the next-hop router. Each next-hop group supports up to 16 next-hop addresses. Up to 30 next-hop groups are supported. Each next-hop group must have at least two next-hop addresses.</p> <p><i>group-name</i>—Name of next-hop group. Up to 30 next-hop groups are supported for the router. Each next-hop group is expected to have at least two next-hop addresses.</p> <p><i>interface-name</i>—Name of interface used to reach the next-hop destination.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Port Mirroring on page 931 |

no-filter-check

| | |
|---------------------------------|--|
| Syntax | no-filter-check; |
| Hierarchy Level | [edit forwarding-options port-mirroring family (inet inet6) output] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>Disable filter checking on the port-mirroring interface.</p> <p>This statement is required when you send port-mirrored traffic to a Tunnel PIC that has a filter applied to it.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Port Mirroring on page 931 |

no-remote-trace (Trace Options)

| | |
|---------------------------------|---|
| Syntax | no-remote-trace; |
| Hierarchy Level | [edit forwarding-options port-mirroring traceoptions],
[edit forwarding-options sampling traceoptions] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Disable remote tracing. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Tracing Traffic Sampling Operations on page 878 |

no-syslog

| | |
|---------------------------------|---|
| Syntax | no-syslog; |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Disable system logging of control protocol requests and responses. By default, these messages are logged. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring System Logging on page 853 |

notification-targets

| | |
|---------------------------------|---|
| Syntax | <code>notification-targets <i>address</i> port <i>port-number</i>;</code> |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | List of destination IP addresses and User Datagram Protocol (UDP) ports to which DFC PICs log exception information and control protocol state transitions, such as timeout values. |
| Options | <code>address <i>address</i></code> —Allowed destination IP address.
<code>port <i>port-number</i></code> —Allowed destination UDP port number. |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Control Source on page 851 |

observation-domain-id

| | |
|---------------------------------|---|
| Syntax | <code>observation-domain-id <i>domain-id</i>;</code> |
| Hierarchy Level | [edit services flow-monitoringversion-ipfix template <i>template-name</i>] |
| Release Information | Statement introduced in Junos OS Release 14.1. |
| Description | <p>For IPFIX flows, an identifier of an Observation Domain is locally unique to an exporting process of the templates. The export process uses the Observation Domain ID to uniquely identify to the collection process in which the flows were metered. We recommend that you configure this ID to be unique for each IPFIX flow. A value of 0 indicates that no specific Observation Domain is identified by this information element. Typically, this attribute is used to limit the scope of other information elements. If the observation domain is not unique, the collector cannot uniquely identify an IPFIX device.</p> <p>If you configure the same Observation Domain ID for different template types, such as for IPv4 and IPv6, it does not impact flow monitoring because the actual or the base observation domain ID is transmitted in the flow. The actual observation domain ID is derived from the value you configure and also in conjunction with other parameters such as the slot number, lookup chip (LU) instance, Packet Forwarding Engine instance. Such a method of computation of the observation domain ID ensures that this ID is not the same for two IPFIX devices.</p> |
| Options | <p><i>domain-id</i>—Specify a unique identifier for the observation domain for IPFIX flows.</p> <p>Range: 0 through 255</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows on page 918• Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows on page 921 |

one-way-hardware-timestamp

| | |
|---------------------------------|---|
| Syntax | one-way-hardware-timestamp; |
| Hierarchy Level | [edit services rpm probe owner test test-name] |
| Release Information | Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.3 for EX Series switches. |
| Description | Enable timestamping of RPM probe messages for one-way delay and jitter measurements. You must configure this statement along with the destination-interface statement to invoke timestamping. This feature is supported only with icmp-ping , icmp-ping-timestamp , udp-ping , and udp-ping-timestamp probe types. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring RPM Timestamping on page 964• destination-interface on page 1652• hardware-timestamp on page 1684 |

option-refresh-rate

| | |
|---------------------------------|--|
| Syntax | option-refresh-rate packets <i>packets</i> seconds <i>seconds</i> ; |
| Hierarchy Level | [edit services flow-monitoring version9],
[edit services flow-monitoring version9 template <i>template-name</i>]
[edit services flow-monitoringversion-ipfix template <i>template-name</i>] |
| Release Information | Statement introduced in Junos OS Release 8.3.
Support at the [edit services flow-monitoring version-ipfix template <i>template-name</i>]
hierarchy level added in Junos OS Release 10.2. |
| Description | Specify the refresh rate, in either packets or seconds. |
| Options | <p><i>packets</i>—Refresh rate, in number of packets.
Range: 1 through 480,000
Default: 4800</p> <p><i>seconds</i>—Refresh rate, in number of seconds.
Range: 10 through 600
Default: 600</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902 |

options-template-id

| | |
|---------------------------------|---|
| Syntax | <code>options-template-id <i>id</i>;</code> |
| Hierarchy Level | <code>[edit services flow-monitoring version9 template <i>template-name</i>]</code>
<code>[edit services flow-monitoringversion-ipfix template <i>template-name</i>]</code> |
| Release Information | Statement introduced in Junos OS Release 14.1. |
| Description | Define a unique options template ID to be used for flow aggregation of version 9 and IPFIX flows. If you do not configure values for the template ID and options template ID, default values are assumed for these IDs, which are different for the various address families. If you configure the same template ID or options template ID value for different address families, such a setting is not processed properly and might cause unexpected behavior. For example, if you configure the same template ID value for both IPv4 and IPv6, the collector validates the export data based on the template ID value that it last receives. In this case, if IPv6 is configured after IPv4, the value is effective for IPv6 and the default value is used for IPv4. |
| Options | <i>id</i> —Specify a unique identifier for the options template to be used for version 9 or IPFIX flows.
Range: 1024 through 65535 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows on page 918 • Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows on page 921 |

output (Accounting)

Syntax `output {
 aggregate-export-interval seconds;
 cflowd hostname {
 aggregation {
 autonomous-system;
 destination-prefix;
 protocol-port;
 source-destination-prefix {
 caida-compliant;
 }
 source-prefix;
 }
 autonomous-system-type (origin | peer);
 (local-dump | no-local-dump);
 port port-number;
 source-address address;
 version format;
 }
 flow-active-timeout seconds;
 flow-inactive-timeout seconds;
 interface interface-name {
 engine-id number;
 engine-type number;
 source-address address;
 }
 }
 }`

Hierarchy Level [edit forwarding-options **accounting** *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure cflowd, output interfaces, and flow properties.

The statements are explained separately.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation • [Configuring Discard Accounting on page 883](#)

output (Monitoring)

Syntax

```
output {
  cflowd hostname port port-number;
  export-format format;
  flow-active-timeout seconds;
  flow-export-destination {
    (cflowd-collector | collector-pic);
  }
  flow-inactive-timeout seconds;
  interface interface-name {
    engine-id number;
    engine-type number;
    input-interface-index number;
    output-interface-index number;
    source-address address;
  }
}
```

Hierarchy Level [edit forwarding-options [monitoring](#) name family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure cflowd, output interfaces, and flow properties.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Flow Monitoring on page 818](#)

output (Port Mirroring)

| | |
|---------------------------------|---|
| Syntax | <pre>output {
 interface <i>interface-name</i> {
 next-hop <i>address</i>;
 }
 no-filter-check;
}</pre> |
| Hierarchy Level | [edit forwarding-options port-mirroring family (inet inet6)] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>Configure output interfaces and flow properties.</p> <p>The statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Port Mirroring on page 931 |

output (Sampling)

```
Syntax  output {
    aggregate-export-interval seconds;
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    extension-service service-name;
    flow-server hostname {
        aggregation {
            autonomous-system;
            destination-prefix;
            protocol-port;
            source-destination-prefix {
                caida-compliant;
            }
            source-prefix;
        }
        autonomous-system-type (origin | peer);
        (local-dump | no-local-dump);
        port port-number;
        source-address address;
        version format;
        version9 {
            template template-name;
        }
    }
    interface interface-name {
        engine-id number;
        engine-type number;
        source-address address;
    }
    file {
        disable;
        filename filename;
        files number;
        size bytes;
        (stamp | no-stamp);
        (world-readable | no-world-readable);
    }
    inline-jflow {
        source-address address;
        flow-export-rate rate;
    }
}
```

Hierarchy Level [edit forwarding-options **sampling instance** *instance-name* **family** (inet | inet6 | mpls)],
[edit forwarding-options **sampling family** (inet | inet6 | mpls)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure cflowd or flow monitoring, output files and interfaces, and flow properties.

The statements are explained separately.



NOTE: The inline-jflow statement is valid only under the [edit forwarding-options sampling instance *instance-name* family inet output] hierarchy level. The file statement is valid only under the [edit forwarding-options sampling family inet output] hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Traffic Sampling on page 871](#)

output-interface-index

Syntax output-interface-index *number*;

Hierarchy Level [edit forwarding-options [monitoring name](#) [output interface interface-name](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify a value for the output interface index that overrides the default supplied by SNMP.

Options *number*—Output interface index value.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Flow Monitoring on page 818](#)

passive-monitor-mode

Syntax passive-monitor-mode;

Hierarchy Level [edit interfaces *interface-name* [unit logical-unit-number](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Asynchronous Transfer Mode (ATM), SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, monitor packet flows from another router. If you include this statement in the configuration, the SONET/SDH interface does not send keepalives or alarms, and does not participate actively on the network.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Enabling Passive Flow Monitoring on page 830](#)
- [multiservice-options on page 1710](#)

password (Flow Collector File Servers)

| | |
|---------------------------------|--|
| Syntax | <code>password "password";</code> |
| Hierarchy Level | [edit services flow-collector destination ftp:url] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the primary and secondary destination FTP server password. |
| Options | <i>password</i> —FTP server password. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Destination FTP Servers for Flow Records on page 840 |

password (Transfer Log File Servers)

| | |
|---------------------------------|---|
| Syntax | <code>password "password";</code> |
| Hierarchy Level | [edit services flow-collector transfer-log-archive archive-sites] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the primary and secondary destination FTP server password. |
| Options | <i>password</i> —FTP server password. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Transfer Logs on page 842 |

peer-as-billing-template

| | |
|---------------------------------|---|
| Syntax | peer-as-billing-template; |
| Hierarchy Level | [edit services flow-monitoring version9 template <i>template-name</i>] |
| Release Information | Statement introduced in Junos OS Release 10.4. |
| Description | Enables the extraction of bandwidth usage information for billing purposes in PIC-based sampling configurations. This capability is supported on routers and applies only to IPv4 and IPv6 traffic. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902 |

pic-memory-threshold

| | |
|---------------------------------|---|
| Syntax | pic-memory-threshold percentage <i>percentage</i> ; |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i>] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Specify a PIC memory usage percentage that triggers a system log warning message. |
| Options | <i>percentage</i> —PIC memory threshold value. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Thresholds on page 854 |

pop-all-labels

| | |
|---------------------------------|--|
| Syntax | pop-all-labels {
required-depth <i>number</i> ;
} |
| Hierarchy Level | [edit interfaces <i>interface-name</i> atm-options mpls],
[edit interfaces <i>interface-name</i> fastether-options mpls],
[edit interfaces <i>interface-name</i> gigether-options mpls],
[edit interfaces <i>interface-name</i> sonet-options mpls] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>For passive monitoring on ATM, SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, removes up to two MPLS labels from incoming IP packets. For passive monitoring on T Series devices, removes up to five MPLS labels from incoming IP packets.</p> <p>This statement has no effect on IP packets with more than two MPLS labels, or IP packets with more than five MPLS labels on T Series devices. Packets with MPLS labels cannot be processed by the monitoring PIC; if packets with MPLS labels are forwarded to the monitoring PIC, they are discarded.</p> <p>The remaining statement is explained separately.</p> |
| Default | If you omit this statement, the MPLS labels are not removed, and the packet is not processed by the monitoring PIC. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Passive Flow Monitoring for MPLS Encapsulated Packets on page 832 • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

port (Flow Monitoring)

| | |
|---------------------------------|--|
| Syntax | <code>port <i>port-number</i>;</code> |
| Hierarchy Level | [edit forwarding-options accounting name output cflowd hostname],
[edit forwarding-options monitoring name family inet output cflowd hostname],
[edit forwarding-options sampling instance instance-name family (inet inet6 mpls) output flow-server hostname],
[edit forwarding-options sampling family (inet inet6 mpls) output flow-server hostname] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the User Datagram Protocol (UDP) port number on the cflowd host system or flow server. |
| Options | <i>port-number</i> —Any valid UDP port number on the host system. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Enabling Flow Aggregation on page 898 |

port (RPM)

| | |
|---------------------------------|---|
| Syntax | <code>port <i>number</i>;</code> |
| Hierarchy Level | [edit services rpm probe-server (tcp udp)] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| Description | Specify the port number for the probe server. |
| Options | <i>number</i> —Port number for the probe server. The value can be 7 or 49,160 through 65,535. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring RPM Receiver Servers on page 963 |

port (TWAMP)

| | |
|---------------------------------|--|
| Syntax | <code>port <i>number</i>;</code> |
| Hierarchy Level | [edit services rpm twamp server] |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | TWAMP server listening port. |
| Options | <i>number</i> —Port number.
Range: 1 through 65,535 |
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring TWAMP on page 968 |

pre-rewrite-tos

| | |
|---------------------------------|--|
| Syntax | <code>pre-rewrite-tos;</code> |
| Hierarchy Level | [edit forwarding-options sampling] |
| Release Information | Statement introduced in Junos OS Release 14.1 |
| Description | Preserve prenormalized type-of-service (ToS) value for egress sampled or mirrored packets. This configuration preserves the prerewrite ToS value for all forms of sampling, such as Routing Engine-based sampling, port mirroring, flow monitoring, and so on. This statement is effective for egress sampling only. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Traffic Sampling on page 871 |

probe

Syntax `probe owner {
 test test-name {
 data-fill data;
 data-size size;
 destination-interface interface-name;
 destination-port port;
 dscp-code-point dscp-bits;
 hardware-timestamp;
 history-size size;
 moving-average-size number;
 one-way-hardware-timestamp;
 probe-count count;
 probe-interval seconds;
 probe-type type;
 routing-instance instance-name;
 source-address address;
 target (url | address);
 test-interval interval;
 thresholds thresholds;
 traps traps;
 }
 }`

Hierarchy Level [edit [services](#) rpm]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify an owner name. The owner name combined with the test name represent a single RPM configuration instance.

Options *owner*—Specify an owner name up to 32 characters in length.

 The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring RPM Probes on page 959](#)

probe-count

| | |
|---------------------------------|--|
| Syntax | <code>probe-count count;</code> |
| Hierarchy Level | <code>[edit services rpm bgp],</code>
<code>[edit services rpm probe owner test test-name]</code>
<code>[edit services rpm twamp client control-connection control-client-name]</code> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.</p> <p>Statement at the <code>[edit services rpm twamp client control-connection control-client-name]</code> hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.</p> |
| Description | Specify the number of probes within a test. |
| Options | <i>count</i> —A value from 1 through 15. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring BGP Neighbor Discovery Through RPM on page 971 • Configuring RPM Probes on page 959 • <i>Two-Way Active Measurement Protocol Overview</i> |

probe-interval

| | |
|---------------------------------|--|
| Syntax | <code>probe-interval <i>interval</i>;</code> |
| Hierarchy Level | [edit services rpm bgp],
[edit services rpm probe owner test <i>test-name</i>]
[edit services rpm twamp client control-connection <i>control-client-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.
Statement at the [edit services rpm twamp client control-connection <i>control-client-name</i>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | Specify the time to wait between sending packets, in seconds. |
| Options | <i>interval</i> —Number of seconds, from 1 through 255. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring BGP Neighbor Discovery Through RPM on page 971• Configuring RPM Probes on page 959• Two-Way Active Measurement Protocol Overview |

probe-limit

| | |
|---------------------------------|---|
| Syntax | <code>probe-limit <i>limit</i>;</code> |
| Hierarchy Level | [edit services rpm] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| Description | Configure the maximum number of concurrent probes allowed. |
| Options | <i>limit</i> —Maximum number of concurrent probes allowed.
Range: 1 through 500(PTX Series Packet Transport Routers only) 1 through 200
Default: 100 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Limiting the Number of Concurrent RPM Probes on page 964 |

probe-server

```
Syntax  probe-server {
        tcp {
            destination-interface interface-name;
            port number;
        }
        udp {
            destination-interface interface-name;
            port number;
        }
    }
```

Hierarchy Level [edit [services](#) rpm]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.

Description Specify the server to act as a receiver for the probes.

The remaining statements are explained separately.



NOTE: The `destination-interface` statement is not supported on PTX Series routers.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring RPM Receiver Servers on page 963](#)

probe-type

| | |
|---------------------------------|--|
| Syntax | <code>probe-type type;</code> |
| Hierarchy Level | <code>[edit services rpm bgp],</code>
<code>[edit services rpm probe owner test test-name]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| Description | Specify the packet and protocol contents of a probe. |
| Options | <p>type—Specify one of the following probe type values:</p> <ul style="list-style-type: none">• http-get—(Not available at the <code>[edit services rpm bgp]</code> hierarchy level.) Sends a Hypertext Transfer Protocol (HTTP) get request to a target URL.• http-metadata-get—(Not available at the <code>[edit services rpm bgp]</code> hierarchy level.) Sends an HTTP get request for metadata to a target URL.• icmp-ping—Sends ICMP echo requests to a target address.• icmp-ping-timestamp—Sends ICMP timestamp requests to a target address.• tcp-ping—Sends TCP packets to a target.• udp-ping—Sends UDP packets to a target.• udp-ping-timestamp—Sends UDP timestamp requests to a target address. |
| Required Privilege Level | <code>interface</code> —To view this statement in the configuration.
<code>interface-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring BGP Neighbor Discovery Through RPM on page 971 |

rate (Forwarding Options)

| | |
|---------------------------------|--|
| Syntax | <code>rate number;</code> |
| Hierarchy Level | [edit forwarding-options analyzer <i>analyzer-name</i> input]
[edit forwarding-options port-mirroring <i>input</i>],
[edit forwarding-options <i>sampling input</i>],
[edit forwarding-options <i>sampling instance instance-name input</i>],
[edit forwarding-options port-mirroring <i>family</i> (inet inet6) <i>input</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
The [edit forwarding-options analyzer <i>analyzer-name</i> input] hierarchy level for MX Series routers introduced in Junos OS Release 14.1. |
| Description | Set a ratio of the number of packets to be sampled. For example, if you specify a rate of 10, every tenth packet (1 packet out of 10) is sampled.

Native analyzer sessions (that is, the [edit forwarding-options analyzer <i>analyzer-name input</i>] hierarchy level for MX Series routers) can be configured without specifying input parameters, which would mean that the instance uses default input values: rate = 1 and maximum-packet-length = 0. |
| Options | <i>number</i> —Denominator of the ratio.
Range: 1 through 65,535 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring Port Mirroring</i> • <i>Configuring Traffic Sampling</i> |


receive-options-packets

| | |
|---------------------------------|---|
| Syntax | <code>receive-options-packets;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> <i>unit</i> <i>logical-unit-number</i> <i>family</i> inet] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | When you enable passive monitoring, this statement is required for conformity with cflowd records structure. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Enabling Passive Flow Monitoring on page 830 |

receive-ttl-exceeded

| | |
|---------------------------------|---|
| Syntax | receive-ttl-exceeded; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | When you enable passive monitoring, this statement is required for conformity with cflowd records structure. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Enabling Passive Flow Monitoring on page 830 |

reflect-mode (RFC2544 Benchmarking)

| | |
|---------------------------------|--|
| Syntax | reflect-mode (mac-rewrite mac-swap no-mac-swap); |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests test-name test-name] |
| Release Information | Statement introduced in Junos OS Release 12.3X52 for ACX Series routers.
Statement introduced in Junos OS Release 14.2 for MX104 3D Universal Edge Routers. |
| Description | Specify the reflection mode for the benchmarking test. |
| Options | <p>mac-rewrite—(ACX Series routers only) Enable rewriting of the MAC address on the reflected frames. The MAC addresses specified in the source-mac-address and destination-mac-address options are used.</p> <p>mac-swap—Swaps the source and destination MAC addresses in the test frame. This is the default behavior.</p> |
| | <div>  <p>NOTE: In bridge families, when the service type is ELAN, MAC addresses are swapped by default, on the reflected frames. And, when the service type is ELINE , MAC addresses are not swapped by default.</p> </div> |
| | <p>no-mac-swap—Does not swap the source and destination MAC addresses in the test frame. The frame is returned to the originator without any modification to the MAC addresses.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Layer 2 RFC2544-Based Benchmarking Tests Overview on page 986 • rfc2544-benchmarking on page 1740 • RFC2544-Based Benchmarking Tests Overview on page 983 • Configuring an RFC 2544-Based Benchmarking Test on page 989 |

required-depth

| | |
|---------------------------------|--|
| Syntax | <code>required-depth <i>number</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> atm-options mpls pop-all-labels],
[edit interfaces <i>interface-name</i> fastether-options mpls pop-all-labels],
[edit interfaces <i>interface-name</i> gigether-options mpls pop-all-labels],
[edit interfaces <i>interface-name</i> sonet-options mpls pop-all-labels] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>For passive monitoring on ATM, SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, specify the number of MPLS labels an incoming packet must have for the pop-all-labels statement to take effect.</p> <p>If you include the required-depth 1 statement, the pop-all-labels statement takes effect for incoming packets with one label only. If you include the required-depth 2 statement, the pop-all-labels statement takes effect for incoming packets with two labels only.</p> |
| Options | <p>number—Number of MPLS labels on incoming IP packets.</p> <p>Range: 1 through 2 labels.</p> <p>Default: If you omit this statement, the pop-all-labels statement takes effect for incoming packets with one or two labels. The default is equivalent to including the required-depth [1 2] statement.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Passive Flow Monitoring for MPLS Encapsulated Packets on page 832• <i>Junos OS Network Interfaces Library for Routing Devices</i> |

retry (Services Flow Collector)

| | |
|---------------------------------|--|
| Syntax | <code>retry number;</code> |
| Hierarchy Level | [edit services flow-collector] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the maximum number of attempts the flow collector interface will make to transfer log files to the FTP server. |
| Options | <i>number</i> —Maximum number of transfer retry attempts.
Range: 0 through 10 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Retry Attempts on page 843 |

retry-delay

| | |
|---------------------------------|---|
| Syntax | <code>retry-delay seconds;</code> |
| Hierarchy Level | [edit services flow-collector] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the amount of time the flow collector interface waits between retry attempts. |
| Options | <i>seconds</i> —Amount of time between transfer retry attempts.
Range: 0 through 60 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Retry Attempts on page 843 |

rfc2544-benchmarking

Syntax rfc2544-benchmarking {
 tests{
 test-name *test-name* {
 test-interface *interface-name*;
 mode reflect;
 family (bridge| inet | ccc);
 destination-ipv4-address *address*;
 destination-udp-port *port-number*;
 source-ipv4-address *address*;
 source-udp-port *port-number*;
 direction (egress | ingress);
 }
 }
 }

Hierarchy Level [edit [services](#) rpm]

Release Information Statement introduced in Junos OS Release 12.3X52 for ACX Series routers.
 Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge Routers.

Description Configure the parameters for the RFC 2544-based benchmarking test. You must configure a test profile, which specifies the type of test and the manner in which it must be performed, and associate the test profile with a test name. The test name that you configure contains details, such as the address family and the test mode, for the test. You can associate the same test profile with multiple test names.


Options **rfc2544-benchmarking**—Define the attributes for the RFC 2544-based benchmarking test to examine and analyze the performance characteristics of a network interconnecting device.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring an RFC 2544-Based Benchmarking Test on page 989](#)
- [RFC2544-Based Benchmarking Tests Overview on page 983](#)
- [show services rpm rfc2544-benchmarking on page 2259](#)
- [show services rpm rfc2544-benchmarking test-id on page 2264](#)

routing-instance

| | |
|---|--|
| Syntax | <code>routing-instance <i>instance-name</i>;</code> |
| Hierarchy Level | <code>[edit services rpm probe owner test <i>test-name</i>]</code>
<code>[edit services rpm twamp client control-connection <i>control-client-name</i>]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.
Statement at the <code>[edit services rpm twamp client control-connection <i>control-client-name</i>]</code> hierarchy level introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | Specify the routing instance used by the probes. |
| <div>  NOTE: Routing instance is also applicable for control connection. </div> | |
| Options | <i>instance-name</i> —A routing instance configured at the <code>[edit routing-instance]</code> hierarchy level.
Default: Internet routing table <code>inet.0</code> . |
| Required Privilege Level | <code>interface</code> —To view this statement in the configuration.
<code>interface-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring RPM Probes on page 959 • <i>Two-Way Active Measurement Protocol Overview</i> |

routing-instances

| | |
|---------------------------------|---|
| Syntax | <code>routing-instances <i>instance-name</i>;</code> |
| Hierarchy Level | [edit services rpm bgp],
[edit services rpm bgp logical-system <i>logical-system-name</i>] |
| Release Information | Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| Description | Specify the routing instance used by the probes. |
| Options | <i>instance-name</i> —A routing instance configured at the [edit routing-instances] hierarchy level.
Default: Internet routing table <code>inet.0</code> . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring BGP Neighbor Discovery Through RPM on page 971 |

rpm (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>rpm (client server twamp-client twamp-server);</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | Associate an RPM client (router or switch that originates RPM probes) or RPM server with a specified interface. |
| Options | <i>client</i> —Identifier for RPM client router or switch.

<i>server</i> —Identifier for RPM server.

<i>twamp-client</i> —Identifier for RPM twamp-client router.

<i>twamp-server</i> —Identifier for RPM twamp-server. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring RPM Timestamping on page 964• Two-Way Active Measurement Protocol Overview |

rpm (Services)

```
Syntax  rpm {
    bgp {
        data-fill data;
        data-size size;
        destination-port port;
        history-size size;
        logical-system logical-system-name [routing-instances routing-instance-name];
        moving-average-size number;
        probe-count count;
        probe-interval seconds;
        probe-type type;
        routing-instances instance-name;
        test-interval interval;
    }
    probe owner {
        test test-name {
            data-fill data;
            data-size size;
            destination-interface interface-name;
            destination-port port;
            dscp-code-point dscp-bits;
            hardware-timestamp;
            history-size size;
            moving-average-size number;
            one-way-hardware-timestamp;
            probe-count count;
            probe-interval seconds;
            probe-type type;
            routing-instance instance-name;
            source-address address;
            target (url url | address address);
            test-interval interval;
            thresholds thresholds;
            traps traps;
        }
    }
    probe-server {
        tcp {
            destination-interface interface-name;
            port number;
        }
        udp {
            destination-interface interface-name;
            port number;
        }
    }
    probe-limit limit;
    traceoptions {
        file filename <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
        flag flag;
    }
}
```

```
twamp {
  server {
    authentication-mode (authenticated | encrypted | none);
    authentication-key-chain identifier {
      key-id identifier {
        secret password-string;
      }
    }
    client-list list-name {
      [ address address ];
    }
    inactivity-timeout seconds;
    maximum-connections-duration hours;
    maximum-connections count;
    maximum-connections-per-client count;
    maximum-sessions count;
    maximum-sessions-per-connection count;
    port number;
    server-inactivity-timeout minutes;
  }
}
rfc2544-benchmarking {
  tests {
    test-name test-name {
      test-interface interface-name;
      mode reflect;
      family (inet | ccc);
      destination-ipv4-address address;
      destination-udp-port port-number;
      source-ipv4-address address;
      source-udp-port port-number;
      direction (egress | ingress);
    }
  }
}
```

| | |
|--------------------------|---|
| Hierarchy Level | [edit services] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure BGP neighbor discovery through RPM.

The remaining statements are explained separately. |
| Usage Guidelines | See “Configuring BGP Neighbor Discovery Through RPM” on page 971 . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |

run-length

| | |
|---------------------------------|--|
| Syntax | <code>run-length <i>number</i>;</code> |
| Hierarchy Level | [edit forwarding-options port-mirroring input],
[edit forwarding-options port-mirroring instance <i>port-mirroring-instance-name</i> input],
[edit forwarding-options port-mirroring family (inet inet6) input],
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance instance-name input] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers. |
| Description | Set the number of samples following the initial trigger event. The configuration enables you to sample packets following those already being sampled. |
| Options | <i>number</i> —Number of samples.
Range: 0 through 20
Default: 0 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Applying Forwarding Table Filters • Configuring Port Mirroring on page 931 • Configuring Traffic Sampling on page 871 |

sample-once

| | |
|---------------------------------|---|
| Syntax | <code>sample-once;</code> |
| Hierarchy Level | [edit forwarding-options sampling] |
| Release Information | Statement introduced in Junos OS Release 9.6. |
| Description | Sample traffic for active monitoring only once. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Traffic Sampling on page 871 |

sampling (Forwarding Options)

```
Syntax  sampling {
        disable;
        sample-once;
        family (inet | inet6 | mpls) {
            disable;
            output {
                aggregate-export-interval seconds;
                extension-service service-name;
                file {
                    disable;
                    filename filename;
                    files number;
                    size bytes;
                    (stamp | no-stamp);
                    (world-readable | no-world-readable);
                }
                flow-active-timeout seconds;
                flow-inactive-timeout seconds;
                flow-server hostname {
                    aggregation {
                        autonomous-system;
                        destination-prefix;
                        protocol-port;
                        source-destination-prefix {
                            caida-compliant;
                        }
                        source-prefix;
                    }
                    autonomous-system-type (origin | peer);
                    (local-dump | no-local-dump);
                    port port-number;
                    source-address address;
                    version format;
                    version9 {
                        template template-name;
                    }
                }
                interface interface-name {
                    engine-id number;
                    engine-type number;
                    source-address address;
                }
            }
        }
    }
    input {
        max-packets-per-second number;
        maximum-packet-length bytes;
        rate number;
        run-length number;
    }
    instance instance-name {
        disable;
```

```

family (inet | inet6 | mpls) {
  disable;
  output {
    aggregate-export-interval seconds;
    extension-service service-name;
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    flow-server hostname {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
      autonomous-system-type (origin | peer);
      (local-dump | no-local-dump);
      port port-number;
      source-address address;
      version format;
      version-ipfix {
        template template-name;
      }
      version9 {
        template template-name;
      }
    }
    inline-jflow {
      source-address address;
      flow-export-rate rate;
    }
    interface interface-name {
      engine-id number;
      engine-type number;
      source-address address;
    }
  }
}
input {
  max-packets-per-second number;
  maximum-packet-length bytes;
  rate number;
  run-length number;
}
pre-rewrite-tos;
traceoptions {
  no-remote-trace;
  file filename <files number> <size bytes> <match expression> <world-readable |
  no-world-readable>;
}
}

```

| | |
|---------------------------------|--|
| Hierarchy Level | [edit forwarding-options] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches. |
| Description | Configure traffic sampling.

The statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Traffic Sampling on page 871• <i>Applying Forwarding Table Filters</i>• <i>Collecting Traffic Sampling Output in the Cisco Systems NetFlow Services Export Version 9 Format</i>• <i>Directing Traffic Sampling Output to a Server Running the cflowd Application</i>• <i>Configuring Port Mirroring</i>• <i>Tracing Traffic-Sampling Operations</i> |

sampling (Interfaces)

| | |
|---------------------------------|---|
| Syntax | sampling <i>direction</i> ; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the direction of traffic to be sampled. |
| Options | <p><i>direction</i> can be one of the following:</p> <p>input—Configure at least one expected ingress point.</p> <p>output—Configure at least one expected egress point.</p> <p>input output—On a single interface, configure at least one expected ingress point and one expect egress point.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Junos OS Services Interfaces Library for Routing Devices</i>• Configuring Flow Monitoring on page 818 |

server

| | |
|---------------------------------|---|
| Syntax | <pre>server { client-list <i>list-name</i> { [address <i>address</i>]; } inactivity-timeout <i>seconds</i>; maximum-connections <i>count</i>; maximum-connections-per-client <i>count</i>; maximum-sessions <i>count</i>; maximum-sessions-per-connection <i>count</i>; port <i>number</i>; }</pre> |
| Hierarchy Level | [edit services rpm twamp] |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | TWAMP server configuration settings. |
| Options | The remaining statements are described separately. |
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring TWAMP on page 968 |


server-inactivity-timeout

| | |
|---------------------------------|--|
| Syntax | server-inactivity-timeout <i>minutes</i> ; |
| Hierarchy Level | [edit services rpm twamp server] |
| Release Information | Statement introduced in Junos OS Release 11.1. |
| Description | The maximum time the Two-Way Active Measurement Protocol (TWAMP) server has to finish the TWAMP control protocol negotiation. |
| Options | <p>minutes—Number of minutes the TWAMP server has to finish the TWAMP control protocol negotiation.</p> <p>Default: 15 minutes</p> <p>Range: 1-30 minutes</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring TWAMP on page 968 |

service-port

| | |
|---------------------------------|---|
| Syntax | <code>service-port <i>port-number</i>;</code> |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Identify the User Datagram Protocol (UDP) port number for control protocol requests. |
| Options | <i>port-number</i> —Port number for control protocol request messages. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring the Control Source on page 851 |

service-type (RFC2544 Benchmarking)

| | |
|--|---|
| Syntax | <code>service-type (elan eline) ;</code> |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests <i>test-name</i> <i>test-name</i>] |
| Release Information | Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.
Statement introduced in Junos OS Release 14.2 for MX104 3D Universal Edge Routers. |
| Description | Mention the service under test. Possible values are elan and eline . This statement is applicable only for the bridge family or when the mode is configured as reflect. When the service type is elan , MAC addresses are swapped by default on the reflected frames. The no-mac-swap is not supported in this service type. When the service type is eline , MAC addresses are not swapped by default in the reflected frames. Use the mac-swap option to swap the addresses. |
| <div> NOTE: When you configure the Layer 2 reflection, you can specify the service type under test as ELINE if you want to simulate an ELINE service using bridge encapsulation.</div> | |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• rfc2544-benchmarking on page 1740• Configuring an RFC 2544-Based Benchmarking Test on page 989• RFC2544-Based Benchmarking Tests Overview on page 983 |

services (RPM)

| | |
|---------------------------------|---|
| Syntax | <code>services rpm { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| Description | Define the service rules to be applied to traffic. |
| Options | <code>rpm</code> —Identifies the RPM set of rules statements. |
| Required Privilege Level | <code>interface</code> —To view this statement in the configuration.
<code>interface-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring BGP Neighbor Discovery Through RPM on page 971 • Configuring RPM Probes on page 959 • Configuring RPM Receiver Servers on page 963 • Limiting the Number of Concurrent RPM Probes on page 964 • Configuring RPM Timestamping on page 964 • Configuring TWAMP on page 968 • Enabling RPM for the Junos OS extension-provider package on page 981 |

shared-key

| | |
|---------------------------------|---|
| Syntax | <code>shared-key value;</code> |
| Hierarchy Level | [edit services dynamic-flow-capture <code>capture-group client-name control-source identifier</code>] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | Configure the authentication key value. |
| Options | <code>value</code> —Secret authentication value shared between a control source and destination. |
| Required Privilege Level | <code>interface</code> —To view this statement in the configuration.
<code>interface-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Control Source on page 851 |

size

| | |
|---------------------------------|--|
| Syntax | <code>size bytes;</code> |
| Hierarchy Level | [edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family (inet inet6 mpls) output file],
[edit forwarding-options sampling traceoptions file] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>Specify the maximum size of each file containing sample or log data. The file size is limited by the number of files to be created and the available hard disk space.</p> <p>When a traffic sampling file named sampling-file reaches the maximum size, it is renamed sampling-file.0. When the sampling-file again reaches its maximum size, sampling-file.0 is renamed sampling-file.1 and sampling-file is renamed sampling-file.0. This renaming scheme continues until the maximum number of traffic sampling files is reached. Then the oldest traffic sampling file is overwritten.</p> |
| Options | <p>bytes—Maximum size of each traffic sampling file or trace log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).</p> <p>Syntax: <i>xk</i> to specify KB, <i>xm</i> to specify MB, or <i>xg</i> to specify GB</p> <p>Range: 10 KB through the maximum file size supported on your router</p> <p>Default: 1 MB for sampling data; 128 KB for log information</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Port Mirroring on page 931• Configuring Traffic Sampling on page 871 |

soft-limit

| | |
|---------------------------------|---|
| Syntax | <code>soft-limit <i>bandwidth</i>;</code> |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>] |
| Release Information | Statement introduced in Junos OS Release 9.2. |
| Description | Specify a bandwidth threshold at which congestion notifications are sent to each control source of the criteria that point to this content destination. If the control source is configured with the syslog statement, a log message will also be generated. |
| Options | <i>bandwidth</i> —Soft limit threshold, in bits per second. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Content Destination on page 850 |

soft-limit-clear

| | |
|---------------------------------|--|
| Syntax | <code>soft-limit-clear <i>bandwidth</i>;</code> |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>] |
| Release Information | Statement introduced in Junos OS Release 9.2. |
| Description | Specify a bandwidth threshold at which the latch set by the soft-limit threshold is cleared. |
| Options | <i>bandwidth</i> —Soft-limit clear threshold, in bits per second. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Content Destination on page 850 • soft-limit on page 1753 |

source-address (Forwarding Options)

| | |
|---------------------------------|---|
| Syntax | <code>source-address <i>address</i>;</code> |
| Hierarchy Level | [edit forwarding-options accounting name output interface <i>interface-name</i>],
[edit forwarding-options monitoring name family <i>family</i> inet output interface <i>interface-name</i>],
[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output interface <i>interface-name</i>],
[edit forwarding-options sampling family (inet inet6 mpls) output interface <i>interface-name</i>],
[edit forwarding-options sampling instance <i>instance-name</i> family inet output inline-jflow] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the source address for monitored packets. |
| Options | <i>address</i> —Interface source address. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Discard Accounting on page 883• Configuring Flow Monitoring on page 818• Configuring Traffic Sampling on page 871 |

source-address (Services)

| | |
|---------------------------------|---|
| Syntax | <code>source-address <i>address</i>;</code> |
| Hierarchy Level | [edit services rpm probe owner test <i>test-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| Description | Specify the source IP address used for probes. If the source IP address is not one of the router's or switch's assigned addresses, the packet will use the outgoing interface's address as its source.

The following addresses cannot be used for the source IP address used for probes: <ul style="list-style-type: none"> • 0.0.0.0 • 127.0.0.0/8 (loopback) • 224.0.0.0/4 (multicast) • 255.255.255.255 (broadcast) |
| Options | <i>address</i> —Valid IP address. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring RPM Probes on page 959 |

source-addresses

| | |
|---------------------------------|---|
| Syntax | <code>source-addresses [<i>addresses</i>];</code> |
| Hierarchy Level | [edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>] |
| Release Information | Statement introduced in Junos OS Release 7.4. |
| Description | List of IP addresses from which the control source can send control protocol requests to the Juniper Networks router. |
| Options | <i>address</i> —Allowed IP source address. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Control Source on page 851 |

source-id

| | |
|---------------------------------|---|
| Syntax | <code>source-id <i>source-id</i>;</code> |
| Hierarchy Level | [edit services flow-monitoring version9 template <i>template-name</i>] |
| Release Information | Statement introduced in Junos OS Release 14.1. |
| Description | For version 9 flows, a 32-bit value that identifies the Exporter Observation Domain is called the source ID. NetFlow collectors use the combination of the source IP address and the source ID field to separate different export streams originating from the same exporter. |
| Options | <i>source-id</i> —Specify a unique identifier for the source for version 9 flows.
Range: 0 through 255 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows on page 918• Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows on page 921 |

source-ipv4-address (RFC 2544 Benchmarking)

| | |
|---------------------------------|---|
| Syntax | <code>source-ipv4-address <i>address</i>;</code> |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests <i>test-name</i> <i>test-name</i>] |
| Release Information | Statement introduced in Junos OS Release 12.3X52 for ACX Series routers.
Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge Routers. |
| Description | Specify the source IPv4 address to be used in generated test frames. This parameter is optional for both ccc and inet families. If you do not configure the source IPv4 address for an inet family, the source address of the interface is used to transmit the test frames. |
| Options | <i>address</i> —Valid IPv4 address.
Default: If you do not configure the source IPv4 address for a ccc family, default value of 192.168.1.10 is used. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring an RFC 2544-Based Benchmarking Test on page 989• RFC2544-Based Benchmarking Tests Overview on page 983• rfc2544-benchmarking on page 1740 |

source-mac-address (RFC2544 Benchmarking)

| | |
|---------------------------------|---|
| Syntax | <code>source-mac-address <i>mac-address</i>;</code> |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests <i>test-name</i> <i>test-name</i>] |
| Release Information | Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.
Statement introduced in Junos OS Release 14.2 for MX104 3D Universal Edge Routers. |
| Description | Specify the source MAC address used in generated test frames. This parameter is applicable for a bridge family. |
| Options | <i>mac-address</i> —Source MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i> ; for example, 0011.2233.4455 or 00:11:22:33:44:55 . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • rfc2544-benchmarking on page 1740 • RFC2544-Based Benchmarking Tests Overview on page 983 • Configuring an RFC 2544-Based Benchmarking Test on page 989 |

source-udp-port (RFC 2544 Benchmarking)

| | |
|---------------------------------|---|
| Syntax | <code>source-udp-port <i>port-number</i>;</code> |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests <i>test-name</i> <i>test-name</i>] |
| Release Information | Statement introduced in Junos OS Release 12.3X52 for ACX Series routers.
Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge Routers. |
| Description | Specify the UDP port of the source to be used in the UDP header for the generated frames. If you do not specify the UDP port, the default value of 4041 is used. |
| Options | <i>port-number</i> —Source UDP port number for the test frames
Default: 4041 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring an RFC 2544-Based Benchmarking Test on page 989 • RFC2544-Based Benchmarking Tests Overview on page 983 • rfc2544-benchmarking on page 1740 |

stamp

| | |
|---------------------------------|---|
| Syntax | (stamp no-stamp); |
| Hierarchy Level | [edit forwarding-options sampling family (inet inet6 mpls) output file] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Include a timestamp with each line in the output file. |
| Options | no-stamp —Do not include timestamps. This is the default.
stamp —Include a timestamp with each line of packet sampling information.
Default: No timestamp is included. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Traffic Sampling on page 871 |

syslog

| | |
|---------------------------------|---|
| Syntax | (syslog no-syslog); |
| Hierarchy Level | [edit interfaces <i>mo-fpc/pic/port</i> multiservice-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | System logging is enabled by default. The system log information of the Monitoring Services PIC is passed to the kernel for logging in the /var/log directory. <ul style="list-style-type: none">• syslog—Enable PIC system logging.• no-syslog—Disable PIC system logging. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Flow Monitoring on page 818 |

target (Services RPM)

| | |
|---------------------------------|---|
| Syntax | <code>target (url <i>url</i> address <i>address</i>);</code> |
| Hierarchy Level | <code>[edit services rpm probe <i>owner</i> test <i>test-name</i>]</code>
<code>[edit services rpm twamp client control-connection <i>control-client-name</i>]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Packet Transport Routers.
Statement at the <code>[edit services rpm twamp client control-connection <i>control-client-name</i>]</code> hierarchy level introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | Specify the destination address or URL used for the probes. |
| Options | url <i>url</i> —For HTTP probe types, specify a fully formed URL that includes http:// in the URL address.

address <i>address</i> —For all other probe types, specify an IPv4 address for the target host. |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring RPM Probes on page 959 • Two-Way Active Measurement Protocol Overview |

tcp

| | |
|---------------------------------|---|
| Syntax | <code>tcp {
 <code>destination-interface</code> <i>interface-name</i>;
 <code>port</code> <i>port</i>;
}</code> |
| Hierarchy Level | <code>[edit <code>services</code> rpm <code>probe-server</code>]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| Description | Specify the port information for the TCP server.

The remaining statements are explained separately. |
| Required Privilege Level | interface —To view this statement in the configuration.
interface-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring RPM Receiver Servers on page 963 |

templates

```
Syntax  templates {
        template-name {
            interval-duration interval-duration;
            inactive-timeout inactive-timeout;
            rate {
                (layer3 layer3-packets-per-second | media media-bits-per-second);
            }
            delay-factor {
                ;
                threshold {
                    (info | warning | critical) delay-factor-threshold;
                }
            }
            media-loss-rate {
                disable;
                threshold {
                    (info | warning | critical) percentage mlr-percentage | packet-count mlr-packet-count;
                }
            }
            media-rate-variation {
                disable;
                threshold {
                    (info | warning | critical) mrp-variation;
                }
            }
            media-packets-count-in-layer3 media-packets-count-in-layer3;
            media-packet-size media-packet-size;
        }
    }
```

Hierarchy Level [edit services [video-monitoring](#)]

Release Information Statement introduced in Junos OS Release 14.1.

Description Configure the media delivery index template containing the measurement parameters for video monitoring.

Options **delay-factor**—Define delay factor syslog threshold levels.

delay-factor-threshold—Delay factor threshold in milliseconds. When the threshold is exceeded, a syslog message is generated.

Default: 0—Do not generate syslogs.

Range: 0 though 65535 milliseconds

disable—Disable logging for the threshold.

inactive-timeout—Number of seconds of flow inactivity after which time media delivery index statistics collection for a flow is terminated.

Range: 30 through 300 seconds

info | warning | critical—Level of syslog message generated when a threshold is exceeded.

interval-duration—Number of seconds after which time media delivery index flow monitoring statistics for the interval are reported.

Range: 1 through 50

layer3-packets-per-second—Layer 3 packet rate in IP packets per second.

Range: 0 through 4,294,967,295 pps

media-bits-per-second—Media bit rate for the stream in bits per second.

media-loss-rate—Define media loss rate syslog threshold levels.

media-packets-count-in-layer-3—Number of media packets in an IP packet.

Range: 1 through 32

media-packet-size—Size of media packet in bits.

Default: 188

Range: 1 through 2048

media-rate-variation—Define delay factor syslog threshold levels.

mlr-packet-count—Media loss rate threshold expressed as the number of packets dropped. When the threshold is exceeded, a syslog message is generated.

mlr-percentage—Media loss rate threshold expressed as the percentage of total packets dropped. When the threshold is exceeded, a syslog message is generated.

Range: 0 through 100

mrv-variation—Media rate variation threshold. The variation is the ratio of actual media rate to the configured media rate, expressed as a percentage.

template-name—Name of the template containing media delivery index measurement criteria. The template can be assigned to an interface.

| | |
|---------------------------------|---|
| Required Privilege Level | interface—To view this statement in the configuration. |
| | interface-control—To add this statement to the configuration. |

| | |
|------------------------------|--|
| Related Documentation | <ul style="list-style-type: none"> • Configuring Inline Video Monitoring on page 1045 |
|------------------------------|--|

test

Syntax `test test-name {
 data-fill data;
 data-size size;
 destination-interface interface-name;
 destination-port port;
 dscp-code-point dscp-bits;
 hardware-timestamp;
 history-size size;
 moving-average-size number;
 one-way-hardware-timestamp;
 probe-count count;
 probe-interval seconds;
 probe-type type;
 routing-instance instance-name;
 source-address address;
 target (url url | address address);
 test-interval interval;
 thresholds thresholds;
 traps traps;
 }`

Hierarchy Level [edit [services](#) rpm [probe](#) owner]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.

Description Specify the range of probes over which the standard deviation, average, and jitter are calculated. The test name combined with the owner name represent a single RPM configuration instance.

Options **test-name**—Specify a test name. The name can be up to 32 characters in length.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring RPM Probes on page 959](#)

tests (RFC 2544 Benchmarking)

Syntax

```
tests {
  test-name test-name {
    test-interface interface-name;
    mode reflect;
    family (inet | ccc);
    destination-ipv4-address address;
    destination-udp-port port-number;
    source-ipv4-address address;
    source-udp-port port-number;
    direction (egress | ingress);
  }
}
```

Hierarchy Level [edit [services](#) rpm [rfc2544-benchmarking](#)]

Release Information Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge Routers.

Description Specify the attributes of the test iteration, such as the address family (type of service, IPv4 or Ethernet), the logical interface, test duration, and test packet size, that are used for a benchmarking test to be run. The test name combined with the test profile represent a single real-time performance monitoring (RPM) configuration instance.

Options **tests**—Define the test iteration for the RFC 2544-based benchmarking test.
The remaining statements are explained separately.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.


Related Documentation

- [Configuring an RFC 2544-Based Benchmarking Test on page 989](#)
- [RFC2544-Based Benchmarking Tests Overview on page 983](#)
- [rfc2544-benchmarking on page 1740](#)

test-interface (RFC 2544 Benchmarking)

| | |
|---------------------------------|--|
| Syntax | <code>test-interface <i>interface-name</i>;</code> |
| Hierarchy Level | [edit services rpm rfc2544-benchmarkingtests <i>test-name</i> <i>test-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge Routers. |
| Description | Specify the logical interface on which the RFC 2544-based benchmarking test is run. If you configure an inet family and the test mode to initiate and terminate test frames on the same device, the interface you configure is not effective. Instead, the test is run on the egress logical interface that is determined using route lookup on the specified destination IPv4 address. If you configure an inet family and the test mode to reflect the frames back on the sender from the other end, the logical interface is used as the interface to enable the reflection service (reflection is performed on the packets entering the specified interface). If you not configure the logical interface for reflection test mode, a lookup is performed on the source IPv4 address to determine the interface that hosts the address. |
| Options | <i>interface-name</i> —Name of the logical interface on which the test needs to be run. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring an RFC 2544-Based Benchmarking Test on page 989• RFC2544-Based Benchmarking Tests Overview on page 983• rfc2544-benchmarking on page 1740 |

test-interval

| | |
|---|---|
| Syntax | <code>test-interval <i>frequency</i>;</code> |
| Hierarchy Level | <code>[edit services rpm bgp],</code>
<code>[edit services rpm probe owner test <i>test-name</i>]</code>
<code>[edit services rpm twamp client control-connection <i>control-client-name</i>]</code> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.</p> <p>Statement at the <code>[edit services rpm twamp client control-connection <i>control-client-name</i>]</code> hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.</p> |
| Description | Specify the time to wait between tests, in seconds. |
| Options | <i>frequency</i> —Number of seconds, from 1 through 86400. |
| <div>  NOTE: For TWAMP the number of seconds are from 1 through 255. </div> | |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring BGP Neighbor Discovery Through RPM on page 971 • Configuring RPM Probes on page 959 • <i>Two-Way Active Measurement Protocol Overview</i> |

test-name (RFC 2544 Benchmarking)

Syntax `test-name test-name {
 test-interface interface-name;
 mode reflect;
 family (inet | ccc);
 destination-ipv4-address address;
 destination-udp-port port-number;
 source-ipv4-address address;
 source-udp-port port-number;
 direction (egress | ingress);
 }`

Hierarchy Level [edit [services](#) rpm [rfc2544-benchmarking tests](#)]

Release Information Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge Routers.

Description Define the name of the RFC 2544-based benchmarking test. For each unique test name that you configure, you can specify a test profile, which contains the settings for a test and its type, and also a test interface, which contains the settings for test packets that are sent and received on the selected interface.

Options *test-name*—Specify a test name. The name can be up to 32 characters in length.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring an RFC 2544-Based Benchmarking Test on page 989](#)
- [RFC2544-Based Benchmarking Tests Overview on page 983](#)
- [rfc2544-benchmarking on page 1740](#)

thresholds

| | |
|---------------------------------|--|
| Syntax | <code>thresholds thresholds;</code> |
| Hierarchy Level | <code>[edit services rpm probe owner test test-name]</code>
<code>[edit services rpm twamp client control-connection control-client-name]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Packet Series Transport Routers.
Statement at the <code>[edit services rpm twamp client control-connection control-client-name]</code> hierarchy level introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | Specify thresholds used for the probes. A system log message is generated when the configured threshold is exceeded. Likewise, an SNMP trap (if configured) is generated when a threshold is exceeded. |
| Options | <p>thresholds—Specify one or more threshold measurements. The following options are supported:</p> <ul style="list-style-type: none"> • egress-time—Measures maximum source-to-destination time per probe. • ingress-time—Measures maximum destination-to-source time per probe. • jitter-egress—Measures maximum source-to-destination jitter per test. • jitter-ingress—Measures maximum destination-to- source jitter per test. • jitter-rtt—Measures maximum jitter per test, from 0 through 60,000,000 microseconds. • rtt—Measures maximum round-trip time per probe, in microseconds. • std-dev-egress—Measures maximum source-to-destination standard deviation per test. • std-dev-ingress—Measures maximum destination-to-source standard deviation per test. • std-dev-rtt—Measures maximum standard deviation per test, in microseconds. • successive-loss—Measures successive probe loss count, indicating probe failure. • total-loss—Measures total probe loss count indicating test failure, from 0 through 15. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring RPM Probes on page 959 • <i>Two-Way Active Measurement Protocol Overview</i> |

traceoptions (Forwarding Options)

| | |
|---------------------------------|---|
| Syntax | <pre>traceoptions {
 no-remote-trace;
 file filename <files number> <size bytes> <match expression> <world-readable
 no-world-readable>;
}</pre> |
| Hierarchy Level | [edit forwarding-options port-mirroring],
[edit forwarding-options sampling] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure traffic sampling tracing operations.

The statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Tracing Traffic Sampling Operations on page 878 |

traceoptions (RPM)

| | |
|----------------------------|--|
| Syntax | <pre> traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i> > <size <i>maximum-file-size</i>> <world-readable no-world-readable>; flag <i>flag</i>; } </pre> |
| Hierarchy Level | [edit services rpm] |
| Release Information | Statement introduced in Junos OS Release 13.2. |
| Description | Define tracing operations for RPM processes. |
| Options | <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. All files are placed in the directory <code>/var/log</code>.</p> <p>Default: <code>rmopd</code></p> <p>files <i>number</i>—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 3 files</p> <p>match <i>regular-expression</i>—(Optional) Refine the output to include lines that contain the regular expression.</p> <p>size <i>maximum-file-size</i>—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the files option.</p> <p>Range: 10 KB through 1 GB</p> <p>Default: 128 KB</p> <p>world-readable—(Optional) Enable unrestricted file access.</p> <p>no-world-readable—(Default) Disable unrestricted file access. This means the log file can be accessed only by the user who configured the tracing operation.</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> • all—Trace all operations. • configuration—Trace configuration events. • error—Trace events related to catastrophic errors in daemon. • ipc—Trace IPC events. • ppm—Trace ppm events. • statistics—Trace statistics. |

Required Privilege Level trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- [Tracing RPM Operations on page 975](#)

transfer

Syntax transfer {
 record-level *number*;
 timeout *seconds*;
}

Hierarchy Level [edit services flow-collector file-specification variant *variant-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify when to send the flow collection file. The file is sent when either of the two conditions is met.

Options record-level *number*—Number of flow collection files collected.

timeout *seconds*—Timeout duration.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring File Formats on page 841](#)

transfer-log-archive

| | |
|---------------------------------|--|
| Syntax | <pre>transfer-log-archive { archive-sites { ftp:url { password "password"; username username; } } filename-prefix prefix; maximum-age minutes; }</pre> |
| Hierarchy Level | [edit services flow-collector] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the filename prefix, maximum age, and destination FTP server for log files containing the transfer activity history for a flow collector interface. |
| Options | The statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Transfer Logs on page 842 |

traps

| | |
|----------------------------|--|
| Syntax | <code>traps traps;</code> |
| Hierarchy Level | [edit services rpm probe owner test test-name]
[edit services rpm twamp client control-connection control-client-name] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.
Statement at the [edit services rpm twamp client control-connection control-client-name] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers. |
| Description | Set the trap bit to generate traps for probes. Traps are sent if the configured threshold is met or exceeded. |
| Options | traps —Specify one or more traps. The following options are supported: <ul style="list-style-type: none">• egress-jitter-exceeded—Generates traps when the jitter in egress time threshold is met or exceeded.• egress-std-dev-exceeded—Generates traps when the egress time standard deviation threshold is met or exceeded.• egress-time-exceeded—Generates traps when the maximum egress time threshold is met or exceeded.• ingress-jitter-exceeded—Generates traps when the jitter in ingress time threshold is met or exceeded.• ingress-std-dev-exceeded—Generates traps when the ingress time standard deviation threshold is met or exceeded.• ingress-time-exceeded—Generates traps when the maximum ingress time threshold is met or exceeded.• jitter-exceeded—Generates traps when the jitter in round-trip time threshold is met or exceeded.• probe-failure—Generates traps for successive probe loss thresholds crossed.• rtt-exceeded—Generates traps when the maximum round-trip time threshold is met or exceeded.• std-dev-exceeded—Generates traps when the round-trip time standard deviation threshold is met or exceeded.• test-completion—Generates traps when a test is completed.• test-failure—Generates traps when the total probe loss threshold is met or exceeded. |



NOTE: For RPM traps to be generated, you must configure the `remote-operations` SNMP trap category by including the `categories` statement at the `[edit snmp trap-group trap-group-name hierarchy level`.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring RPM Probes on page 959](#)
- *categories*
- *Two-Way Active Measurement Protocol Overview*

tth

Syntax `tth hops;`

Hierarchy Level `[edit services dynamic-flow-capture capture-group client-name content-destination identifier]`

Release Information Statement introduced in Junos OS Release 7.4.

Description Time-to-live (TTL) value for the IP-IP header.

Options *hops*—TTL value.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Content Destination on page 850](#)

twamp

Syntax

```
twamp {
  server {
    authentication-mode mode;
    authentication-key-chain identifier {
      key-id identifier {
        secret password-string;
      }
    }
    client-list list-name {
      [ address address ];
    }
    inactivity-timeout seconds;
    max-connection-duration hours;
    maximum-connections count;
    maximum-connections-per-client count;
    maximum-sessions count;
    maximum-sessions-per-connection count;
    port number;
    routing-instance-list {
      instance-name {
        port number;
      }
    }
    server-inactivity-timeout minutes;
  }
}
```

Hierarchy Level [edit services rpm]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the Two-Way Active Measurement Protocol (TWAMP) responder or sever settings on all M Series and T Series routers that support Multiservices PICs (running in either Layer 2 or Layer 3 mode), and on MX Series routers.

TWAMP is an open protocol for measurement of two-way metrics. The host that initiates the TCP connection takes the roles of the control-client and (in the two-host implementation) the session-sender. Such a device is also called the TWAMP client. The host that acknowledges the TCP connection accepts the roles of a server and (in the two-host implementation) and the session-reflector. Such a device is also called the TWAMP server. The TWAMP-Test messages are exchanged between the session-sender and the session-reflector, and the TWAMP-Control messages are exchanged between the control-client and the server.

The following addresses cannot be used for the **client-list** source IP address used for probes:

- 0.0.0.0
- 127.0.0.0/8 (loopback)
- 224.0.0.0/4 (multicast)

- 255.255.255.255 (broadcast)

The remaining statements are described separately.

| | |
|---------------------------------|--|
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring TWAMP on page 968 |

twamp-server

| | |
|---------------------------------|--|
| Syntax | twamp-server; |
| Hierarchy Level | [edit interfaces <i>sp-fpc/pic/port</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced in Junos OS Release 9.3. |
| Description | Specify the service PIC logical interface to provide the TWAMP service. |
| Required Privilege Level | system—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring TWAMP on page 968 |

template (Forwarding Options)

| | |
|---------------------------------|---|
| Syntax | template <i>template-name</i> ; |
| Hierarchy Level | [edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output flow-server <i>hostname version9</i>],
[edit forwarding-options sampling family (inet inet6 mpls) output flow-server <i>hostname version9</i>] |
| Release Information | Statement introduced in Junos OS Release 8.3. |
| Description | Specify flow monitoring version 9 template to be used for output of sampling records. |
| Options | <i>template-name</i> —Name of the version 9 template. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902 |

template-id

| | |
|---------------------------------|--|
| Syntax | template-id <i>id</i> ; |
| Hierarchy Level | [edit services flow-monitoring version9 template <i>template-name</i>]
[edit services flow-monitoringversion-ipfix template <i>template-name</i>] |
| Release Information | Statement introduced in Junos OS Release 14.1. |
| Description | Define a template ID to be used for flow aggregation of version 9 and IPFIX flows. If you do not configure values for the template ID and options template ID, default values are assumed for these IDs, which are different for the various address families. If you configure the same template ID or options template ID value for different address families, such a setting is not processed properly and might cause unexpected behavior. For example, if you configure the same template ID value for both IPv4 and IPv6, the collector validates the export data based on the template ID value that it last receives. In this case, if IPv6 is configured after IPv4, the value is effective for IPv6 and the default value is used for IPv4. |
| Options | <i>id</i> —Specify a unique identifier for the template to be used for version 9 or IPFIX flows.
Range: 1024 through 65535 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows on page 918• Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows on page 921 |

template-refresh-rate

| | |
|---------------------------------|--|
| Syntax | template-refresh-rate packets <i>packets</i> seconds <i>seconds</i> ; |
| Hierarchy Level | [edit services flow-monitoring version9 template <i>template-name</i>]
[edit services flow-monitoringversion-ipfix template <i>template-name</i>] |
| Release Information | Statement introduced in Junos OS Release 8.3.
Support at the [edit services flow-monitoring version-ipfix template <i>template-name</i>] hierarchy level added in Junos OS Release 10.2. |
| Description | Specify the refresh rate, in either packets or seconds. |
| Options | <i>packets</i> —Refresh rate, in number of packets.
Range: 1 through 480,000
Default: 4800

<i>seconds</i> —Refresh rate, in number of seconds.
Range: 10 through 600
Default: 600 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902 |

trio-flow-offload

| | |
|---------------------------------|---|
| Syntax | trio-flow-offload minimum-bytes <i>minimum-bytes</i> ; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> services-options] |
| Release Information | Statement introduced in Junos OS Release 12.1. |
| Description | Enable any plug-in or daemon on a PIC to generate a flow offload request to of-load flows to the Packet Forwarding Engine. This command is available on MX Series routers with Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs). |
| Options | <i>minimum-bytes</i> —The minimum number of bytes that trigger offloading. When this option is omitted, offloading is triggered when both the forward and reverse flows of the session have begun, meaning that at least one packet has flowed in each direction. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Flow Offloading on page 827 |

udp

Syntax `udp {
 destination-interface interface-name;
 port port;
 }`

Hierarchy Level [edit `services` rpm `probe-server`]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.3 for EX Series switches.
 Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.

Description Specify the port information for the UDP server.

 The remaining statements are explained separately.



NOTE: The `destination-interface` statement is not supported on PTX Series routers.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Configuring RPM Receiver Servers on page 963](#)

unit

Syntax

```
unit logical-unit-number {
    family inet {
        address address {
            destination destination-address;
        }
        filter {
            group filter-group-number;
            input filter-name;
            output filter-name;
        }
        sampling direction;
    }
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options *logical-unit-number*—Number of the logical unit.

Range: 0 through 16,384

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Network Interfaces Library for Routing Devices* for other statements that do not affect services interfaces.
- *Junos OS Network Interfaces Library for Routing Devices*

username (Services)

| | |
|---------------------------------|---|
| Syntax | <code>username <i>user-name</i>;</code> |
| Hierarchy Level | [edit services flow-collector transfer-log-archive archive-sites] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the username for the transfer log server. |
| Options | <i>username</i> —FTP server username. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Transfer Logs on page 842 |

variant

| | |
|---------------------------------|---|
| Syntax | <pre>variant <i>variant-number</i> {
 data-format <i>format</i>;
 name-format <i>format</i>;
 transfer {
 record-level <i>number</i>;
 timeout <i>seconds</i>;
 }
}</pre> |
| Hierarchy Level | [edit services flow-collector file-specification] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure a variant of the file format. |
| Options | The statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring File Formats on page 841 |

version

| | |
|---------------------------------|--|
| Syntax | <code>version <i>format</i>;</code> |
| Hierarchy Level | [edit forwarding-options accounting name output flow-server hostname],
[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output flow-server hostname],
[edit forwarding-options sampling family (inet inet6 mpls) output flow-server hostname] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the version format of the aggregated flows exported to a cflowd server. |
| Options | <i>format</i> —Format of the flows.
Values: 5 or 8
Default: 5 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • export-format on page 1663 • Enabling Flow Aggregation on page 898 |

version9 (Forwarding Options)

| | |
|---------------------------------|---|
| Syntax | <code>version9 {
 template <i>template-name</i>;
}</code> |
| Hierarchy Level | [edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output flow-server hostname],
[edit forwarding-options sampling family (inet inet6 mpls) output flow-server hostname] |
| Release Information | Statement introduced in Junos OS Release 8.3. |
| Description | Specify flow monitoring version 9 properties to apply to output sampling records. The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Flow Aggregation to Use Version 9 Flow Templates on page 902 |

video-monitoring

```
Syntax video-monitoring {
  templates {
    template-name {
      interval-duration interval-duration;
      inactive-timeout inactive-timeout;
      rate {
        (layer3 layer3-packets-per-second | media media-bits-per-second);
      }
      delay-factor {
        disable;
        threshold {
          (info | warning | critical) delay-factor-threshold;
        }
      }
      media-loss-rate {
        disable;
        threshold {
          (info | warning | critical) percentage mlr-percentage | packet-count
            mlr-packet-count;
        }
      }
      media-rate-variation {
        ;
        threshold {
          (info | warning | critical) mrp-variation;
        }
      }
      media-packets-count-in-layer3 media-packets-count-in-layer3;
      media-packet-size media-packet-size;
    }
  }
  interfaces {
    interface-name {
      family {
        inet {
          input-flows {
            input-flow-name {
              source-address [ address ];
              destination-address [ address ];
              source-port [ port ];
              destination-port [ port ];
              template template-name;
            }
          }
          output-flows {
            output-flow-name {
              source-address [ address ];
              destination-address [ address ];
              source-port [ port ];
              destination-port [ port ];
              template template-name;
            }
          }
        }
      }
    }
  }
}
```



```

    }
  }
}

```

| | |
|---------------------------------|--|
| Hierarchy Level | [edit services] |
| Release Information | Statement introduced in Junos OS Release 14.1. |
| Description | Define the options for video monitoring using media delivery index options for metrics. The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Inline Video Monitoring on page 1045 |

world-readable

| | |
|---------------------------------|---|
| Syntax | (world-readable no-world-readable); |
| Hierarchy Level | [edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family (inet inet6 mpls) output file],
[edit forwarding-options sampling traceoptionsfile] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Enable unrestricted file access. |
| Options | no-world-readable —Restrict file access to owner. This is the default.

world-readable —Enable unrestricted file access.

Default: no-world-readable |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Port Mirroring on page 931 • Configuring Traffic Sampling on page 871 |

Tunnel and Encryption Services Configuration Statements

- [address \(Interfaces\) on page 1785](#)
- [allow-fragmentation on page 1785](#)
- [backup-destination on page 1786](#)
- [backup-interface on page 1786](#)

- [copy-tos-to-outer-ip-header](#) on page 1787
- [destination \(Interfaces\)](#) on page 1788
- [destination \(Routing Instance\)](#) on page 1789
- [destination \(Tunnel Remote End\)](#) on page 1789
- [destination-networks](#) on page 1790
- [do-not-fragment](#) on page 1791
- [dynamic-tunnels](#) on page 1792
- [es-options](#) on page 1793
- [family](#) on page 1794
- [filter](#) on page 1795
- [hold-time \(OAM\)](#) on page 1795
- [interfaces](#) on page 1796
- [ipsec-sa](#) on page 1796
- [keepalive-time](#) on page 1797
- [key](#) on page 1798
- [multicast-only](#) on page 1798
- [peer-unit](#) on page 1799
- [reassemble-packets](#) on page 1799
- [redundancy-group \(Interfaces\)](#) on page 1800
- [redundancy-group \(Logical Tunnels\)](#) on page 1801
- [routing-instance](#) on page 1802
- [routing-instances](#) on page 1802
- [routing-options](#) on page 1803
- [source](#) on page 1803
- [source](#) on page 1804
- [source-address](#) on page 1804
- [ttl](#) on page 1805
- [tunnel](#) on page 1806
- [tunnel](#) on page 1807
- [unit \(Interfaces\)](#) on page 1808
- [unit \(Interfaces\)](#) on page 1809

address (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <code>address <i>address</i> {
 <i>destination address</i>;
}</code> |
| Hierarchy Level | [edit <code>interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i></code>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the interface address. |
| Options | <p><i>address</i>—Address of the interface.</p> <p>The remaining statement is explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Encryption Interfaces on page 1251 |

allow-fragmentation

| | |
|---------------------------------|---|
| Syntax | <code>allow-fragmentation;</code> |
| Hierarchy Level | <p>[edit <code>interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel</code>],</p> <p>[edit <code>logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel</code>]</p> |
| Release Information | Statement introduced in Junos OS Release 9.2. |
| Description | Enable fragmentation of generic routing encapsulation (GRE) encapsulated packets regardless of maximum transmission unit (MTU) value. |
| Default | By default, the GRE-encapsulated packets are dropped if the packet size exceeds the MTU setting of the egress interface. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • reassemble-packets on page 1799 • Configuring Packet Reassembly on page 1217 |

backup-destination

| | |
|---------------------------------|---|
| Syntax | <code>backup-destination <i>destination-address</i>;</code> |
| Hierarchy Level | <code>[edit interfaces <i>interface-name</i> unit logical-unit-number tunnel],[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit logical-unit-number tunnel]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For tunnel interfaces, specify the remote address of the backup tunnel. |
| Options | <i>destination-address</i> —Address of the remote side of the connection. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• destination (Interfaces) on page 1600• destination (Tunnel Remote End) on page 1789• Configuring IPsec Tunnel Redundancy on page 1260 |

backup-interface

| | |
|---------------------------------|--|
| Syntax | <code>backup-interface <i>interface-name</i>;</code> |
| Hierarchy Level | <code>[edit interfaces <i>interface-name</i> es-options]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure a backup ES Physical Interface Card (PIC). When the primary ES PIC has a servicing failure, the backup becomes active, inherits all the tunnels and security associations (SAs), and acts as the new next hop for IPsec traffic. |
| Options | <i>interface-name</i> —Name of ES interface to serve as the backup. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring ES PIC Redundancy on page 1259 |

copy-tos-to-outer-ip-header

| | |
|---------------------------------|--|
| Syntax | copy-tos-to-outer-ip-header; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced in Junos OS Release 8.2. |
| Description | For GRE tunnel interfaces only, enable the inner IP header's ToS bits to be copied to the outer IP packet header. |
| Default | If you omit this statement, the ToS bits in the outer IP header are set to 0. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header on page 1217 |

destination (Interfaces)

| | |
|---------------------------------|---|
| Syntax | <code>destination address;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],
[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>For CoS on ATM interfaces, specify the remote address of the connection.</p> <p>For point-to-point interfaces only, specify the address of the interface at the remote end of the connection.</p> <p>For tunnel and encryption interfaces, specify the remote address of the tunnel.</p> |
| Options | <i>address</i> —Address of the remote side of the connection. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Linear RED Profiles on ATM Interfaces• Multilink and Link Services Logical Interface Configuration Overview on page 717• Configuring Encryption Interfaces on page 1251• Configuring Traffic Sampling on page 871• Configuring Flow Monitoring on page 818• Configuring Unicast Tunnels on page 1213 |

destination (Routing Instance)

| | |
|---------------------------------|---|
| Syntax | <code>destination <i>routing-instance-name</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel <i>routing-instance</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the destination routing instance that points to the routing table containing the tunnel destination address. |
| Default | The default Internet routing table inet.0 . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Tunnel Interfaces for Routing Table Lookup on page 1241 |

destination (Tunnel Remote End)

| | |
|---------------------------------|---|
| Syntax | <code>destination <i>address</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches. |
| Description | For tunnel interfaces, specify the remote address of the tunnel. |
| Options | <i>destination-address</i> —Address of the remote side of the connection. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Unicast Tunnels on page 1213 • Configuring Traffic Sampling on page 871 • Configuring Flow Monitoring on page 818 |

destination-networks

| | |
|---------------------------------|---|
| Syntax | <code>destination-networks <i>prefix</i>;</code> |
| Hierarchy Level | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>],</code>
<code>[edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry</i>],</code>
<code>[edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>],</code>
<code>[edit routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>],</code>
<code>[edit routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry</i>],</code>
<code>[edit routing-options dynamic-tunnels <i>tunnel-name</i>],</code>
<code>[edit routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry</i>]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3 for ACX Series routers. |
| Description | Specify the IPv4 prefix range for the destination network. Only tunnels within the specified IPv4 prefix range can be created. |
| Options | <i>prefix</i> —Destination prefix of the network. |
| Required Privilege Level | <code>routing</code> —To view this statement in the configuration.
<code>routing-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring GRE Tunnels for Layer 3 VPNs• Configuring Dynamic Tunnels on page 1245• Configuring RSVP Automatic Mesh |

do-not-fragment

| | |
|---------------------------------|--|
| Syntax | do-not-fragment; |
| Hierarchy Level | [edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel],
[edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel] |
| Release Information | Statement introduced in Junos OS Release 9.2. |
| Description | Set the do-not-fragment (DF) bit on the packets entering the GRE tunnel so that they do not get fragmented anywhere in the path. |
| Default | By default, fragmentation is disabled. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• reassemble-packets on page 1799• Configuring Packet Reassembly on page 1217 |

dynamic-tunnels

| | |
|--------------------------|--|
| Syntax | <pre>dynamic-tunnels <i>tunnel-name</i> {
 <i>destination-networks</i> <i>prefix</i>;
 gre;
 rsvp-te <i>entry-name</i> {
 <i>destination-networks</i> <i>network-prefix</i>;
 label-switched-path-template {
 default-template;
 <i>template-name</i>;
 }
 }
 source-address <i>address</i>;
}</pre> |
| Hierarchy Level | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> <i>routing-options</i>],
[edit logical-systems <i>logical-system-name</i> <i>routing-options</i>],
[edit routing-instances <i>routing-instance-name</i> <i>routing-options</i>],
[edit <i>routing-options</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3 for ACX Series routers. |
| Description | Configure a dynamic tunnel between two PE routers. |
| Options | <i>tunnel-name</i> —Name of the dynamic tunnel.

The remaining statements are explained separately. |
| Required Privilege Level | routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks</i>• <i>Configuring GRE Tunnels for Layer 3 VPNs</i>• Configuring Dynamic Tunnels on page 1245 |

es-options

| | |
|---------------------------------|--|
| Syntax | <pre>es-options {
 backup-interface <i>interface-name</i>;
}</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | <p>On ES interfaces, configure ES interface-specific interface properties.</p> <p>The backup-interface statement is explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring ES PIC Redundancy on page 1259 |

family

| | |
|---------------------------------|---|
| Syntax | <pre>family inet {
 ipsec-sa sa-name;
}</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure protocol family information for the logical interface. |
| Options | <p>family—Protocol family:</p> <ul style="list-style-type: none">• ccc—Circuit cross-connect protocol suite• inet—IP version 4 suite• inet6—IP version 6 suite• iso—Open Systems Interconnection (OSI) International Organization for Standardization (ISO) protocol suite• mlfr-end-to-end—Multilink Frame Relay FRF.15• mlfr-uni-nni—Multilink Frame Relay FRF.16• multilink-ppp—Multilink Point-to-Point Protocol• mpls—MPLS• tcc—Translational cross-connect protocol suite• tnp—Trivial Network Protocol• vpls—Virtual private LAN service <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Encryption Interfaces on page 1251• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces. |

filter

| | |
|---------------------------------|--|
| Syntax | filter {
input <i>filter-name</i> ;
output <i>filter-name</i> ;
} |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Define the filters to be applied on an interface. |
| Options | <p>input <i>filter-name</i>—Identifier for the input filter.</p> <p>output <i>filter-name</i>—Identifier for the output filter.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Filters for Traffic Transiting the ES PIC on page 1253 |

hold-time (OAM)

| | |
|---------------------------------|--|
| Syntax | hold-time <i>seconds</i> ; |
| Hierarchy Level | <p>[edit protocols oam],</p> <p>[edit protocols oam gre-tunnel interface <i>interface-name</i>]</p> |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Length of time the originating end of a GRE tunnel waits for keepalive packets from the other end of the tunnel before marking the tunnel as operationally down. |
| Options | <p><i>seconds</i>—Hold-time value.</p> <p>Default: 5 seconds</p> <p>Range: 5 through 250 seconds</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • GRE Keepalive Time Overview on page 1205 • Configuring GRE Keepalive Time on page 1205 • keepalive-time on page 1797 |


interfaces

| | |
|---------------------------------|---|
| Syntax | <code>interfaces { ... }</code> |
| Hierarchy Level | [edit] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure interfaces on the router. |
| Default | The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Junos OS Network Interfaces Library for Routing Devices</i> |

ipsec-sa

| | |
|---------------------------------|---|
| Syntax | <code>ipsec-sa <i>sa-name</i>;</code> |
| Hierarchy Level | [edit interfaces <i>es-fpc/pic/port</i> unit <i>logical-unit-number</i> family inet] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the IP Security (IPsec) SA name associated with the interface. |
| Options | <i>sa-name</i> —IPsec SA name. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Encryption Interfaces on page 1251• <i>Junos OS Administration Library for Routing Devices</i> |

keepalive-time

| | |
|---|---|
| Syntax | keepalive-time <i>seconds</i> ; |
| Hierarchy Level | [edit protocols oam],
[edit protocols oam gre-tunnel interface <i>interface-name</i>],
[edit protocols oam gre-tunnel interface <i>interface-name.unit-number</i>] |
| Release Information | Statement introduced in Junos OS Release 10.2. |
| Description | Time difference between consecutive keepalive packets in a GRE tunnel. |
| <div>  NOTE: Support for GRE keepalive packets on MPC line cards became available as of Junos OS Release 11.4. </div> | |
| Options | <i>seconds</i> —Keepalive time value.
Default: 1 second
Range: 1 through 50 seconds |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • GRE Keepalive Time Overview on page 1205 • Configuring GRE Keepalive Time on page 1205 • hold-time on page 1795 |

key

| | |
|---------------------------------|--|
| Syntax | <code>key number;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Identify an individual traffic flow within a tunnel, as defined in RFC 2890, <i>Key and Sequence Number Extensions to GRE</i> . On M Series and T Series routers, you can configure the GRE interface on an Adaptive Services, Multiservices, or Tunnel PIC. On MX Series routers, configure the interface on a Multiservices DPC. |
| Options | number —Value of the key.
Range: 0 through 4,294,967,295 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring a Key Number on GRE Tunnels on page 1215 |

multicast-only

| | |
|---------------------------------|--|
| Syntax | <code>multicast-only;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure the unit and family so that the interface can transmit and receive multicast traffic only. You can configure this property on the IP family only. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Restricting Tunnels to Multicast Traffic on page 1219• tunnel on page 1807 |

peer-unit

| | |
|---------------------------------|--|
| Syntax | <code>peer-unit <i>unit-number</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure a peer relationship between two logical systems. |
| Options | <i>unit-number</i> —Peering logical system unit number. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Logical Tunnel Interfaces on page 1221 |

reassemble-packets

| | |
|---------------------------------|--|
| Syntax | <code>reassemble-packets;</code> |
| Hierarchy Level | [edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i>],
[edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced in Junos OS Release 9.2. |
| Description | Enable reassembly of fragmented tunnel packets on generic routing encapsulation (GRE) tunnel interfaces. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Packet Reassembly on page 1217 |

redundancy-group (Interfaces)

| | |
|---------------------------------|--|
| Syntax | <pre>redundancy-group {
 member-interface <i>interface-name</i> {
 (active backup);
 }
}</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.3. |
| Description | Configure member logical tunnels of redundant logical tunnels only on MX Series 3D Universal Edge Routers. |
| Options | <p>active—Set the interface to the active mode.</p> <p>backup—Set the interface to the backup mode.</p> <p>The remaining statement is explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To view this statement in the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Example: Configuring Redundant Logical Tunnels on page 1227• Configuring Redundant Logical Tunnels on page 1226• Redundant Logical Tunnels Overview on page 1224• redundancy-group (Logical Tunnels) on page 1801 |

redundancy-group (Logical Tunnels)

| | |
|---------------------------------|--|
| Syntax | <pre> redundancy-group { interface-type { redundant-logical-tunnel { device <i>count</i>; } } } </pre> |
| Hierarchy Level | [edit chassis] |
| Release Information | Statement introduced in Junos OS Release 13.3.
Support for up to 255 redundant logical tunnels added to Junos OS Release 13.3R3. |
| Description | Configure redundant logical tunnels only on MX Series 3D Universal Edge Routers. |
| Options | <p><i>count</i>—Specify the number of the redundant logical tunnels. For Junos OS Release 13.3R1, 13.3R2, and 14.1R1 the valid range is from 1 to 16. For Junos OS Release 13.3R3 the valid range is from 1 to 255.</p> <p>—</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To view this statement in the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Example: Configuring Redundant Logical Tunnels on page 1227 • Configuring Redundant Logical Tunnels on page 1226 • Redundant Logical Tunnels Overview on page 1224 • redundancy-group (Interfaces) on page 1800 |

routing-instance

| | |
|---------------------------------|--|
| Syntax | <code>routing-instance {
 destination <i>routing-instance-name</i>;
}</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel],
[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Specify the destination routing instance that points to the routing table containing the tunnel destination address. |
| Default | The default Internet routing table inet.0 . |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Tunnel Interfaces for Routing Table Lookup on page 1241 |

routing-instances

| | |
|---------------------------------|--|
| Syntax | <code>routing-instances <i>routing-instance-name</i> { ... }</code> |
| Hierarchy Level | [edit],
[edit logical-systems <i>logical-system-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure an additional routing entity for a router or switch. You can create multiple instances of BGP, IS-IS, OSPF, OSPF version 3 (OSPFv3), and RIP for a router or switch. |
| Default | Routing instances are disabled for the router or switch. |
| Options | <i>routing-instance-name</i> —Name of the routing instance, a maximum of 31 characters. The remaining statements are explained separately. |
| Required Privilege Level | routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring EVPN Routing Instances• Configuring EVPN Routing Instances on EX9200 Switches• Configuring Routing Instances on PE Routers in VPNs |

routing-options

| | |
|---------------------------------|--|
| Syntax | routing-options { ... } |
| Hierarchy Level | [edit],
[edit logical-systems <i>logical-system-name</i>],
[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>],
[edit routing-instances <i>routing-instance-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Configure protocol-independent routing properties. |
| Required Privilege Level | routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Protocol-Independent Routing Properties Feature Guide for Routing Devices</i> |

source

| | |
|---------------------------------|--|
| Syntax | source <i>source-address</i> ; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],
[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | For tunnel and encryption interfaces, specify the source address. |
| Options | <i>source-address</i> —Address of the source side of the connection. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Encryption Interfaces on page 1251 • Configuring Traffic Sampling on page 871 • Configuring Flow Monitoring on page 818 |

source

| | |
|---------------------------------|---|
| Syntax | <code>source source-address;</code> |
| Hierarchy Level | <code>[edit interfaces interface-name unit logical-unit-number tunnel]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify the source address of the tunnel. |
| Default | If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel. |
| Options | source-address —Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field. |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Generic Routing Encapsulation Tunneling (CLI Procedure) |

source-address

| | |
|---------------------------------|---|
| Syntax | <code>source-address address;</code> |
| Hierarchy Level | <code>[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options dynamic-tunnels tunnel-name],</code>
<code>[edit logical-systems logical-system-name routing-options dynamic-tunnels tunnel-name],</code>
<code>[edit routing-instances routing-instance-name routing-options dynamic-tunnels tunnel-name],</code>
<code>[edit routing-options dynamic-tunnels tunnel-name]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3 for ACX Series routers. |
| Description | Configure the tunnel source address. |
| Options | address —Name of the source address. |
| Required Privilege Level | routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Dynamic Tunnels on page 1245 |

ttl

| | |
|---------------------------------|---|
| Syntax | <code>ttl value;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>number</i> tunnel] |
| Release Information | Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches. |
| Description | Set the time-to-live value bit in the header of the outer IP packet. |
| Options | value —Time-to-live value.
Range: 0 through 255
Default: 64 |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Tunnel Properties</i>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i> |

tunnel

Syntax tunnel {
 `backup-destination` *destination-address*;
 `destination` *destination-address*;
 `routing-instance` {
 `destination` *routing-instance-name*;
 }
 `source` *source-address*;
 `ttl` *number*;
 }

Hierarchy Level [edit interfaces *interface-name* `unit` *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic or virtual private networks (VPNs).

The statements are explained separately.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related • [Configuring Encryption Interfaces on page 1251](#)
Documentation • *Tunnel Properties*
 • *Junos OS VPNs Library for Routing Devices*

tunnel

| | |
|---------------------------------|---|
| Syntax | <pre> tunnel { allow-fragmentation; backup-destination address; destination destination-address; do-not-fragment; key number; routing-instance { destination routing-instance-name; } source source-address; ttl number; } </pre> |
| Hierarchy Level | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> |
| Description | <p>Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic or virtual private networks (VPNs).</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Encryption Interfaces on page 1251 • <i>Junos OS VPNs Library for Routing Devices</i> |

unit (Interfaces)

Syntax `unit logical-unit-number {
 family inet {
 ipsec-sa sa-name;
 }
 tunnel {
 backup-destination destination-address;
 destination destination-address;
 routing-instance {
 destination routing-instance-name;
 }
 source source-address;
 ttl number;
 }
 }`

Hierarchy Level [edit [interfaces](#) interface-name]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options *logical-unit-number*—Number of the logical unit.

Range: 0 through 16,384

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Encryption Interfaces on page 1251](#)
- *Junos OS Network Interfaces Library for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices* for other statements that do not affect services interfaces.

unit (Interfaces)

| | |
|---------------------------------|---|
| Syntax | <pre> unit logical-unit-number { peer-unit unit-number; reassemble-packets; tunnel { allow-fragmentation; backup-destination address; destination destination-address; do-not-fragment; key number; routing-instance { destination routing-instance-name; } source source-address; ttl number; } } </pre> |
| Hierarchy Level | [edit interfaces interface-name],
[edit logical-systems logical-system-name interfaces interface-name] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device. |
| Options | <p>logical-unit-number—Number of the logical unit.</p> <p>Range: 0 through 16,384</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces. |

Operational Commands

- [Adaptive Services Operational Commands on page 1811](#)
- [Application Aware Services Operational Commands on page 2081](#)
- [Link and Multilink Services Operational Commands on page 2117](#)
- [Monitoring, Sampling, and Collection Services Operational Commands on page 2162](#)
- [Tunnel and Encryption Services Operational Commands on page 2294](#)

Adaptive Services Operational Commands

- `clear services cos statistics`
- `clear services crtp statistics`
- `clear services ids`
- `clear services ids destination-table`
- `clear services ids pair-table`
- `clear services ids source-table`
- `clear services inline nat pool`
- `clear services inline nat statistics`
- `clear services ipsec-vpn certificates`
- `clear services ipsec-vpn ike security-associations`
- `clear services ipsec-vpn ipsec security-associations`
- `clear services ipsec-vpn ipsec statistics`
- `clear services l2tp destination`
- `clear services l2tp destination statistics`
- `clear services l2tp multilink`
- `clear services l2tp session`
- `clear services l2tp session statistics`
- `clear services l2tp tunnel`
- `clear services l2tp tunnel statistics`
- `clear services nat flows`
- `clear services nat mappings`

- clear services nat mappings app
- clear services nat mappings eim
- clear services nat mappings pcp
- clear security pki ca-certificate
- clear security pki certificate-request
- clear security pki crl
- clear security pki key-pair
- clear security pki local-certificate
- clear services service-sets statistics integrity-drops
- clear services service-sets statistics packet-drops
- clear services service-sets statistics syslog
- clear services stateful-firewall flows
- clear services stateful-firewall sip-call
- clear services stateful-firewall sip-register
- clear services stateful-firewall statistics
- request interface (revert | switchover) (Adaptive Services)
- request security pki ca-certificate enroll
- request security pki ca-certificate load
- request security pki ca-certificate verify
- request security pki crl load
- request security pki generate-certificate-request
- request security pki generate-key-pair
- request security pki local-certificate enroll
- request security pki local-certificate generate-self-signed
- request security pki local-certificate load
- request security pki local-certificate verify
- request services ipsec-vpn ipsec switch tunnel
- show interfaces (Adaptive Services)
- show interfaces (Link Services IQ)
- show interfaces (Redundant Adaptive Services)
- show interfaces (Redundant Link Services IQ)
- show interfaces load-balancing
- show interfaces redundancy
- show security pki ca-certificate
- show security pki certificate-request
- show security pki crl
- show security pki local-certificate

- `show services cos statistics`
- `show services crtp`
- `show services crtp flows`
- `show services ids`
- `show services inline nat pool`
- `show services inline nat statistics`
- `show services ipsec-vpn certificates`
- `show services ipsec-vpn ike security-associations`
- `show services ipsec-vpn ipsec security-associations`
- `show services ipsec-vpn ipsec statistics`
- `show services link-services cpu-usage`
- `show services l2tp multilink`
- `show services l2tp radius`
- `show services l2tp session`
- `show services l2tp summary`
- `show services l2tp tunnel`
- `show services l2tp user`
- `show services nat ipv6-multicast-interfaces`
- `show services nat mappings`
- `show services nat pool`
- `show services pcp statistics`
- `show services service-sets cpu-usage`
- `show services service-sets memory-usage`
- `show services service-sets statistics packet-drops`
- `show services service-sets statistics syslog`
- `show services service-sets statistics tcp-mss`
- `show services service-sets summary`
- `show services software`
- `show services software flows`
- `show services software statistics`
- `show services stateful-firewall conversations`
- `show services stateful-firewall flow-analysis`
- `show services stateful-firewall flows`
- `show services stateful-firewall sip-call`
- `show services stateful-firewall sip-register`
- `show services stateful-firewall statistics`

- `show services stateful-firewall statistics application-protocol sip`
- `show services stateful-firewall subscriber-analysis`

clear services cos statistics

| | |
|---------------------------------|---|
| Syntax | <code>clear services cos statistics</code>
<code><interface <i>interface-name</i>></code>
<code><service-set <i>service-set-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 8.1. |
| Description | Clear statistics for class-of-service (CoS) code point bit patterns and forwarding classes as configured in CoS services for the AS PIC. |
| Options | none —Clear all services CoS statistics.

interface <i>interface-name</i> —(Optional) Clear statistics for the specified interface only.

service-set <i>service-set-name</i> —(Optional) Clear statistics for the specified service set only. |
| Required Privilege Level | view |
| List of Sample Output | clear services cos statistics on page 1815 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services cos statistics

```
user@host> clear services cos statistics
```

clear services crtp statistics

| | |
|---------------------------------|--|
| Syntax | clear services crtp statistics
<interface <i>interface-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear Compressed Real-Time Transport Protocol (CRTP) flow statistics. |
| Options | none —Clear CRTP flow statistics on all interfaces.

interface <i>interface-name</i> —(Optional) Clear CRTP flow statistics for the specified interface.
On M Series and T Series routers, a link services IQ (lsq-<i>fpc/pic/port</i>) or redundant link services IQ (rlsq-<i>fpc/pic/port</i>) interface. |
| Required Privilege Level | view |
| List of Sample Output | clear services crtp statistics on page 1816 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services crtp statistics

```
user@host> clear services crtp statistics
```

clear services ids

| | |
|---------------------------------|---|
| Syntax | clear services ids
<interface <i>interface-name</i> >
<service-set <i>service-set-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear intrusion detection service (IDS) events. |
| Options | <p>none—Clear all IDS events for all adaptive services interfaces for all service sets, and clear and reset IDS.</p> <p>interface <i>interface-name</i>—(Optional) On M Series and T Series routers, the <i>interface-name</i> can be <i>sp-fpc/pic/port</i> or <i>rspnumber</i>.</p> <p>service-set <i>service-set-name</i>—(Optional) Clear all IDS events for a particular service set.</p> |
| Required Privilege Level | view |
| List of Sample Output | clear services ids on page 1817 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services ids

```
user@host> clear services ids
```

clear services ids destination-table

| | |
|---------------------------------|--|
| Syntax | <code>clear services ids destination-table</code>
<code><destination-prefix <i>destination-prefix-name</i>></code>
<code><interface <i>interface-name</i>></code>
<code><service-set <i>service-set-name</i>></code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear the intrusion detection service (IDS) events for a particular address that might be under attack. |
| Options | <p>none—Clear the attack destination address table.</p> <p>destination-prefix <i>destination-prefix-name</i>—(Optional) Clear the attack destination table for a particular destination prefix.</p> <p>interface <i>interface-name</i>—(Optional) Clear the attack destination table for a particular interface. On M Series and T Series routers, the <i>interface-name</i> can be <i>sp-fpc/pic/port</i> or <i>rspnumber</i>.</p> <p>service-set <i>service-set-name</i>—(Optional) Clear the attack destination table for a particular service set.</p> |
| Required Privilege Level | view |
| List of Sample Output | clear services ids destination-table on page 1818 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services ids destination-table

```
user@host> clear services ids destination-table
```

clear services ids pair-table

| | |
|---------------------------------|--|
| Syntax | clear services ids pair-table
<destination-prefix <i>destination-prefix-name</i> >
<interface <i>interface-name</i> >
<service-set <i>service-set-name</i> >
<source-prefix <i>source-prefix-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear the intrusion detection service (IDS) attack source and destination address pair table. |
| Options | <p>none—Clear the attack source and destination address pair table.</p> <p>destination-prefix <i>destination-prefix-name</i>—(Optional) Clear the attack source and destination address pair table for a particular destination prefix.</p> <p>interface <i>interface-name</i>—(Optional) Clear the attack destination table for a particular interface. On M Series and T Series routers, the <i>interface-name</i> can be sp-fpc/pic/port or rspnumber.</p> <p>service-set <i>service-set-name</i>—(Optional) Clear the attack source and destination address pair table for a particular service set.</p> <p>source-prefix <i>source-prefix-name</i>—(Optional) Clear the attack source and destination address pair table for a particular source prefix.</p> |
| Required Privilege Level | view |
| List of Sample Output | clear services ids pair-table on page 1819 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services ids pair-table

```
user@host> clear services ids pair-table
```

clear services ids source-table

| | |
|---------------------------------|---|
| Syntax | <code>clear services ids source-table</code>
<code><interface <i>interface-name</i>></code>
<code><service-set <i>service-set-name</i>></code>
<code><source-prefix <i>source-prefix-name</i>></code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear all intrusion detection service (IDS) events for addresses that are suspected attackers. |
| Options | <p>none—Clear the attack source address table.</p> <p>interface <i>interface-name</i>—(Optional) On M Series and T Series routers, the <i>interface-name</i> can be <i>sp-fpc/pic/port</i> or <i>rspnumber</i>.</p> <p>service-set <i>service-set-name</i>—(Optional) Clear the attack source address table for a particular service set.</p> <p>source-prefix <i>source-prefix-name</i>—(Optional) Clear the attack source address table for a particular source prefix.</p> |
| Required Privilege Level | view |
| List of Sample Output | clear services ids source-table on page 1820 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services ids source-table

```
user@host> clear services ids source-table
```

clear services inline nat pool

| | |
|---------------------------------|--|
| Syntax | clear services inline nat pool <i>pool-name</i> |
| Release Information | Command introduced in Junos OS Release 11.4. |
| Description | Clear global inline NAT statistics. |
| Options | pool-name —Name of the NAT pool for which statistic are cleared. |
| Required Privilege Level | clear |
| List of Sample Output | clear services inline nat pool on page 1821 |
| Output Fields | When you enter this command, the NAT pool statistics are cleared. There is no specific output. |

Sample Output

clear services inline nat pool

```
user@host> clear services inline nat pool p1
```

clear services inline nat statistics

| | |
|---------------------------------|--|
| Syntax | clear services inline nat statistics
<interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 11.4. |
| Description | Clear global inline NAT statistics. |
| Options | interface <i>interface-name</i> —(Optional) Clear inline NAT statistics for the specified interface only. |
| Required Privilege Level | clear |
| List of Sample Output | clear services inline nat statistics on page 1822 |
| Output Fields | When you enter this command, the global inline NAT statistics are cleared. There is no specific output. |

Sample Output

clear services inline nat statistics

```
user@host> clear services inline nat statistics
```


clear services ipsec-vpn certificates

| | |
|---------------------------------|---|
| Syntax | clear services ipsec-vpn certificates (all service-set <i>service-set</i>)
<certificate-cache-entry <i>number</i> > |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | (Adaptive services interfaces only) Delete digital certificates from the IPsec configuration memory cache. Issuing this command also clears the certificate revocation list (CRL) from the cache along with the certificates. |
| Options | all —Delete digital certificates for all service sets.

service-set <i>service-set</i> —Delete digital certificates for the specified service set. |
| Required Privilege Level | clear |
| List of Sample Output | clear services ipsec-vpn certificates all on page 1823 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services ipsec-vpn certificates all

```
user@host> clear services ipsec-vpn certificates all
```

clear services ipsec-vpn ike security-associations

| | |
|---------------------------------|---|
| Syntax | clear services ipsec-vpn ike security-associations
<peer-address-name>
<service-set service-set-name> |
| Release Information | Command introduced before Junos OS Release 7.4.
service-set option added in Junos OS Release 8.5. |
| Description | (Adaptive services interfaces only) Clear Internet Key Exchange (IKE) security associations. |
| Options | peer-address-name —(Optional) Clear only the security association specified by the peer address.

service-set service-set-name —(Optional) Clear only the security association specified by the service-set name. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show services ipsec-vpn ike security-associations on page 1961 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services ipsec-vpn ike security-associations

```
user@host> clear services ipsec-vpn ike security-associations
```

clear services ipsec-vpn ipsec security-associations

| | |
|---------------------------------|--|
| Syntax | clear services ipsec-vpn security-associations
<peer-address-name>
<remote-gateway remote-gateway-address>
<service-set-name>
<tunnel-index tunnel-index-number> |
| Release Information | Command introduced before Junos OS Release 7.4.
remote-gateway , service-set-name , and tunnel-index options added in Junos OS Release 8.4. |
| Description | (Adaptive services interfaces only) Clear IP Security (IPsec) security associations. You can combine the options for greater specificity. |
| Options | <p>peer-address-name—(Optional) Clear only the security association specified by the peer address.</p> <p>remote-gateway remote-gateway-address—(Optional) Clear only the security association specified by the remote gateway address.</p> <p>service-set-name—(Optional) Clear only the security association specified by the service-set name.</p> <p>tunnel-index tunnel-index-number—(Optional) Clear only the security association specified by the tunnel index number.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show services ipsec-vpn ipsec security-associations on page 1965 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services ipsec-vpn ipsec security-associations

```
user@host> clear services ipsec-vpn ipsec security-associations
```

clear services ipsec-vpn ipsec statistics

| | |
|---------------------------------|--|
| Syntax | clear services ipsec-vpn ipsec statistics
<remote-gateway <i>address</i> >
<service-set <i>service-set-name</i> > |
| Release Information | Command introduced in Junos OS Release 8.1. |
| Description | (Adaptive services interface only) Clear IP Security (IPsec) statistics. |
| Options | remote-gateway <i>address</i> —(Optional) Clear statistics for the specified remote system.
service-set <i>service-set-name</i> —(Optional) Clear statistics for the specified service set. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show services ipsec-vpn ipsec statistics on page 1969 |
| List of Sample Output | clear services ipsec-vpn ipsec statistics on page 1826 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services ipsec-vpn ipsec statistics

```
user@host> clear services ipsec-vpn ipsec statistics
```

clear services l2tp destination

| | |
|---------------------------------|---|
| Syntax | clear services l2tp destination
<all local-gateway <i>gateway-address</i> peer-gateway <i>gateway-address</i> > |
| Release Information | Command introduced in Junos OS Release 10.4.
Statistics option introduced in Junos OS Release 13.1 |
| Description | Clear all Layer 2 Tunneling Protocol (L2TP) destinations and all tunnels and sessions that belong to the destinations. This command is available only for LAC on MX Series routers. |
| Options | <p>all—Close all L2TP destinations.</p> <p>local-gateway <i>gateway-address</i>—Clear only the L2TP destinations and all tunnels and sessions associated with the specified local gateway address.</p> <p>peer-gateway <i>gateway-address</i>—Clear only the L2TP destinations and all tunnels and sessions associated with the peer gateway with the specified address.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show services l2tp destination |
| List of Sample Output | clear services l2tp destination all on page 1827 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services l2tp destination all

```
user@host> clear services l2tp destination all

Destination 2 closed
```

clear services l2tp destination statistics

| | |
|---------------------------------|---|
| Syntax | <code>clear services l2tp destination statistics</code>
<code><all local-gateway <i>gateway-address</i> peer-gateway <i>gateway-address</i> ></code> |
| Release Information | Command introduced in Junos OS Release 13.1. |
| Description | Clear all statistics associated with the Layer 2 Tunneling Protocol (L2TP) destinations and all tunnels and sessions that belong to the destinations. This command is available only for LAC on MX Series routers. |
| Options | <p>all—Clear all statistics associated with the L2TP destinations.</p> <p>local-gateway <i>gateway-address</i>—Clear statistics related to L2TP destination and all tunnels and sessions associated with the specified local gateway address.</p> <p>peer-gateway <i>gateway-address</i>—Clear statistics related to L2TP destination and all tunnels and sessions associated with the specified peer gateway address.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• <i>show services l2tp destination</i> |
| List of Sample Output | clear services l2tp destination statistics on page 1828 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services l2tp destination statistics

```
user@host>clear services l2tp destination statistics all
Destination 1 statistics cleared
```

clear services l2tp multilink

| | |
|---------------------------------|---|
| Syntax | clear services l2tp multilink (all <statistics> bundle-id <i>number</i> <statistics> statistics (all bundle-id <i>number</i>)) |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M10i and M7i routers only) Close Layer 2 Tunneling Protocol (L2TP) multilink sessions or clear session statistics. |
| Options | <p>all <statistics>—Close all L2TP multilink sessions or clear statistics for all L2TP multilink sessions.</p> <p>bundle-id <i>number</i> <statistics>—L2TP multilink bundle ID. The value is an internally generated number from 1 to 65535. Close the specified L2TP multilink session, or using the statistics keyword with this option, clear statistics for the specified session.</p> <p>statistics (all bundle-id <i>number</i>)—Clear all session statistics or clear statistics for the specified multilink bundle ID.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • L2TP Services Configuration Overview on page 638 • L2TP Minimum Configuration on page 639 • show services l2tp multilink on page 1976 |
| List of Sample Output | clear services l2tp multilink statistics all on page 1829 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services l2tp multilink statistics all

```
user@host> clear services l2tp multilink statistics all
Multilink 1 statistics cleared
```

clear services l2tp session

| | |
|---------------------------------|--|
| Syntax | <code>clear services l2tp session (all interface <i>interface-name</i> local-gateway <i>gateway-address</i> local-gateway-name <i>gateway-name</i> local-session-id <i>session-id</i> local-tunnel-id <i>tunnel-id</i> peer-gateway <i>gateway-address</i> peer-gateway-name <i>gateway-name</i> tunnel-group <i>group-name</i> user <i>username</i>)</code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M10i and M7i routers only) Clear Layer 2 Tunneling Protocol (L2TP) sessions on LNS.

(MX Series routers only) Clear L2TP sessions on LAC and LNS. |
| Options | <p>all—Close all L2TP sessions.</p> <p>interface <i>interface-name</i>—Clear only the L2TP sessions using the specified adaptive services or inline services interface. The interface type depends on the line card as follows:</p> <ul style="list-style-type: none">• si-<i>fpc/pic/port</i>—MPCs on MX Series routers only. This option is not available for L2TP on M Series routers.• sp-<i>fpc/pic/port</i>—AS or Multiservices PICs on M7i, M10i, and M120 routers only. This option is not available for L2TP on MX Series routers. <p>local-gateway <i>gateway-address</i>—Clear only the L2TP sessions associated with the specified local gateway address.</p> <p>local-gateway-name <i>gateway-name</i>—Clear only the L2TP sessions associated with the specified local gateway name.</p> <p>local-session-id <i>session-id</i>—Clear only the L2TP sessions with this identifier for the local endpoint of the L2TP session.</p> <p>local-tunnel-id <i>tunnel-id</i>—Clear only the L2TP sessions associated with the specified local tunnel identifier.</p> <p>peer-gateway <i>gateway-address</i>—Clear only the L2TP sessions associated with the peer gateway with the specified address.</p> <p>peer-gateway-name <i>gateway-name</i>—Clear only the L2TP sessions associated with the peer gateway with the specified name.</p> <p>tunnel-group <i>group-name</i>—Clear only the L2TP sessions associated with the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.</p> <p>user <i>username</i>—(M Series routers only) Clear only the L2TP sessions for the specified username.</p> |
| Required Privilege Level | clear |

- Related Documentation**
- [L2TP Services Configuration Overview on page 638](#)
 - [L2TP Minimum Configuration on page 639](#)
 - [clear services l2tp session statistics on page 1832](#)
 - [show services l2tp session on page 1984](#)

List of Sample Output [clear services l2tp session on page 1831](#)
[clear services l2tp session interface on page 1831](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services l2tp session

```
user@host> clear services l2tp session 31694
```

```
Session 31694 closed
```

Sample Output

clear services l2tp session interface

```
user@host> show services l2tp session Tunnel local ID: 17185
```

| Local ID | Remote ID | State | Interface unit | Interface Name |
|----------|-----------|-------------|----------------|----------------|
| 5117 | 1 | Established | 1073741828 | si-2/0/0 |
| 34915 | 2 | Established | 1073741829 | si-2/1/0 |
| 6454 | 3 | Established | 1073741830 | si-2/0/0 |
| 46142 | 4 | Established | 1073741831 | si-2/1/0 |

```
user@host> clear services l2tp session interface si-2/0/0
```

```
Session 5117 closed
```

```
Session 6454 closed
```

```
user@host> show services l2tp session Tunnel local ID: 17185
```

| Local ID | Remote ID | State | Interface unit | Interface Name |
|----------|-----------|-------------|----------------|----------------|
| 34915 | 2 | Established | 1073741829 | si-2/1/0 |
| 46142 | 4 | Established | 1073741831 | si-2/1/0 |

clear services l2tp session statistics

| | |
|---------------------------------|--|
| Syntax | <code>clear services l2tp session statistics (all interface <i>interface-name</i> local-gateway <i>gateway-address</i> local-gateway-name <i>gateway-name</i> local-session-id <i>session-id</i> local-tunnel-id <i>tunnel-id</i> peer-gateway <i>gateway-address</i> peer-gateway-name <i>gateway-name</i> tunnel-group <i>group-name</i> user <i>username</i>)</code> |
| Release Information | Command introduced before Junos OS Release 7.4.
Support for MX Series routers added in Junos OS Release 10.4. |
| Description | (M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Clear statistics for Layer 2 Tunneling Protocol (L2TP) sessions. |
| Options | <p>all—Clear statistics for all L2TP sessions.</p> <p>interface <i>interface-name</i>—Clear only the L2TP sessions using the specified adaptive services or inline services interface. The interface type depends on the line card as follows:</p> <ul style="list-style-type: none">• si-<i>fpc/pic/port</i>—MPCs on MX Series routers only. This option is not available for L2TP on M Series routers.• sp-<i>fpc/pic/port</i>—AS or Multiservices PICs on M7i, M10i, and M120 routers only. This option is not available for L2TP on MX Series routers. <p>local-gateway <i>gateway-address</i>—Clear statistics for only the L2TP sessions associated with the local gateway with the specified address.</p> <p>local-gateway-name <i>gateway-name</i>—Clear statistics for only the L2TP sessions associated with the local gateway with the specified name.</p> <p>local-session-id <i>session-id</i>—Clear statistics for only the L2TP sessions with this identifier for the local endpoint of the L2TP session.</p> <p>local-tunnel-id <i>tunnel-id</i>—Clear statistics for only the L2TP sessions associated with the specified local tunnel identifier.</p> <p>peer-gateway <i>gateway-address</i>—Clear statistics for only the L2TP sessions associated with the peer gateway with the specified address.</p> <p>peer-gateway-name <i>gateway-name</i>—Clear statistics for only the L2TP sessions associated with the peer gateway with the specified name.</p> <p>tunnel-group <i>group-name</i>—Clear statistics for only the L2TP sessions associated with the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.</p> <p>user <i>username</i> <statistics>—Clear statistics for only the L2TP sessions for the specified username. This option is not available for L2TP LAC on MX Series routers.</p> |
| Required Privilege Level | view |

- Related Documentation**
- [L2TP Services Configuration Overview on page 638](#)
 - [L2TP Minimum Configuration on page 639](#)
 - [clear services l2tp session on page 1830](#)
 - [show services l2tp session on page 1984](#)

List of Sample Output [clear services l2tp session statistics all on page 1833](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[clear services l2tp session statistics all](#)

```
user@host> clear services l2tp session statistics all
Session 26497 statistics cleared
```

clear services l2tp tunnel

| | |
|---------------------------------|---|
| Syntax | clear services l2tp tunnel (all interface <i>sp-fpc/pic/port</i> local-gateway <i>gateway-address</i> local-gateway-name <i>gateway-name</i> local-tunnel-id <i>tunnel-id</i> peer-gateway <i>gateway-address</i> peer-gateway-name <i>gateway-name</i> tunnel-group <i>group-name</i>) |
| Release Information | Command introduced before Junos OS Release 7.4.
Support for LAC on MX Series routers introduced in Junos OS Release 10.4.
Support for LNS on MX Series routers introduced in Junos OS Release 11.4. |
| Description | (M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Clear Layer 2 Tunneling Protocol (L2TP) tunnels. |
| Options | <p>all—Clear all L2TP tunnels.</p> <p>sp-fpc/pic/port—(Optional) Clear only the L2TP tunnels using the specified adaptive services interface. This option is not available for L2TP on MX Series routers.</p> <p>local-gateway gateway-address—Clear only the L2TP tunnels associated with the local gateway with the specified address.</p> <p>local-gateway-name gateway-name—Clear only the L2TP tunnels associated with the local gateway with the specified name.</p> <p>local-tunnel-id tunnel-id—Clear only the L2TP tunnels that have the specified local tunnel identifier.</p> <p>peer-gateway gateway-address—Clear only the L2TP tunnels associated with the peer gateway with the specified address.</p> <p>peer-gateway-name gateway-name—Clear only the L2TP tunnels associated with the peer gateway with the specified name.</p> <p>tunnel-group group-name—Clear only the L2TP tunnels in the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• L2TP Services Configuration Overview on page 638• L2TP Minimum Configuration on page 639• clear services l2tp tunnel statistics on page 1836• show services l2tp tunnel on page 1997 |
| List of Sample Output | clear services l2tp tunnel on page 1835 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services l2tp tunnel

```
user@host> clear services l2tp tunnel 17185
```

```
Tunnel 17185 closed
```

clear services l2tp tunnel statistics

| | |
|---------------------------------|---|
| Syntax | <code>clear services l2tp tunnel statistics (all interface <i>sp-fpc/pic/port</i> local-gateway <i>gateway-address</i> local-gateway-name <i>gateway-name</i> local-tunnel-id <i>tunnel-id</i> peer-gateway <i>gateway-address</i> peer-gateway-name <i>gateway-name</i> tunnel-group <i>group-name</i>)</code> |
| Release Information | Command introduced before Junos OS Release 7.4.
Support for MX Series routers added in Junos OS Release 10.4. |
| Description | (M10i and M7i routers: LNS only. MX Series routers: LAC only.) Clear statistics for Layer 2 Tunneling Protocol (L2TP) tunnels. |
| Options | <p>all—Clear statistics for all L2TP tunnels.</p> <p>interface <i>sp-fpc/pic/port</i>—Clear statistics for only the L2TP tunnels using the specified adaptive services interface. This option is not available for L2TP LAC on MX Series routers.</p> <p>local-gateway <i>gateway-address</i>—Clear statistics for only the L2TP tunnels associated with the local gateway with the specified address.</p> <p>local-gateway-name <i>gateway-name</i>—Clear statistics for only the L2TP tunnels associated with the local gateway with the specified name.</p> <p>local-tunnel-id <i>tunnel-id</i>—Clear statistics for only the L2TP tunnels that have the specified local tunnel identifier.</p> <p>peer-gateway <i>gateway-address</i>—Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified address.</p> <p>peer-gateway-name <i>gateway-name</i>—Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified name.</p> <p>tunnel-group <i>group-name</i>—Clear statistics for only the L2TP tunnels in the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• L2TP Services Configuration Overview on page 638• L2TP Minimum Configuration on page 639• clear services l2tp tunnel on page 1834• show services l2tp tunnel on page 1997 |
| List of Sample Output | clear services l2tp tunnel statistics all on page 1837 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

`clear services l2tp tunnel statistics all`

```
user@host> clear services l2tp tunnel statistics all
Tunnel 9933 statistics cleared
```

clear services nat flows

| | |
|---------------------------------|--|
| Syntax | clear services nat flows
<b4address b4address>
<service-set service-set>
<subscriber subscriber-address> |
| Release Information | Command introduced in Junos OS Release 14.1. |
| Description | Clear NAT flows. |
| Options | <p>none—Clear all NAT flows.</p> <p>b4address b4address—(Optional) Clear NAT flows for a particular B4 address.</p> <p>service-set service-set—(Optional) Clear NAT flows for a particular service set.</p> <p>subscriber ip—(Optional) Clear NAT flows for a particular subscriber, identified by IPv4 address.</p> |
| Required Privilege Level | view |
| Related Documentation | |
| List of Sample Output | clear services nat flows subscriber (IPv4 address) on page 1838 |
| Output Fields | Table 54 on page 1838 lists the output fields for the clear services nat flows command. Output fields are listed in the approximate order in which they appear. |

Table 54: clear services nat flows Output Fields

| Field Name | Field Description |
|----------------------|---|
| Interface | Name of a services interface. |
| Service set | Name of the service set from which flows are being cleared. |
| Flows removed | Number of flows removed. |

Sample Output

clear services nat flows subscriber (IPv4 address)

```

user@host> clear services nat flows subscriber ip 3.3.3.3
Interface  Service set  Flows removed

sp-2/0/0   ss1             0

```

Sample Output

clear services nat mappings

| | |
|---------------------------------|--|
| Syntax | clear services nat mappings
<app>
<eim>
<pcp>
<service-set <i>service-set</i> > |
| Release Information | Command introduced in Junos OS Release 14.1. |
| Description | Clear NAT mappings. |
| Options | <p>none—Clear all NAT mappings.</p> <p>app—(Optional) Clear address-pooling paired NAT mappings.</p> <p>eim—(Optional) Clear endpoint-independent NAT mappings.</p> <p>pcp—(Optional) Clear Port Control Protocol NAT mappings.</p> <p>service-set <i>service-set</i>—(Optional) Clear NAT mappings for a specified service set..</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show services nat mappings on page 2009 • clear services nat mappings app on page 1841 • clear services nat mappings eim on page 1842 • clear services nat mappings pcp on page 1844 |
| List of Sample Output | clear services nat mappings on page 1840 |
| Output Fields | <p>Table 55 on page 1839 lists the output fields for the clear services nat mappings command. Output fields are listed in the approximate order in which they appear.</p> |

Table 55: clear services nat mappings Output Fields

| Field Name | Field Description |
|-------------------------|---|
| Interface | Name of a services interface. |
| Service set | Name of the service set from which flows are being cleared. |
| Mappings removed | Number of mappings removed. |
| Flows removed | Number of flows removed. |

Sample Output

clear services nat mappings

```
user@host> clear services nat mappings
Interface  Service set  Mappings removed  Flows removed
sp-2/0/0   ss1          0                  0
```

clear services nat mappings app

| | |
|---------------------------------|---|
| Syntax | clear services nat mappings app
<b4address <i>b4address/prefix</i> >
<service-set <i>service-set</i> >
<subscriber <i>subscriber-ipv4-address</i> > |
| Release Information | Command introduced in Junos OS Release 14.1. |
| Description | Clear NAT mappings for address pooling paired (app). |
| Options | <p>none—Clear all NAT app mappings.</p> <p>b4address <i>b4address/prefix</i>—(Optional) Clear NAT APP mappings for a particular subscriber <i>b4address/prefix</i></p> <p>service-set <i>service-set</i>—(Optional) Clear NAT APP mappings for a specified service set..</p> <p>subscriber <i>subscriber-ipv4-address/prefix</i>—(Optional) Clear NAT APP mappings for a particular subscriber <i>ipv4-address/prefix</i></p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show services nat mappings on page 2009 |
| List of Sample Output | clear services nat mappings app on page 1841 |
| Output Fields | Table 56 on page 1841 lists the output fields for the clear services nat mappings app command. Output fields are listed in the approximate order in which they appear. |

Table 56: clear services nat mappings app Output Fields

| Field Name | Field Description |
|-------------------------|---|
| Interface | Name of a services interface. |
| Service set | Name of the service set from which flows are being cleared. |
| Mappings removed | Number of mappings removed. |
| Flows removed | Number of flows removed. |

Sample Output

clear services nat mappings app

```

user@host> clear services nat mappings app
Interface  Service set  Mappings removed  Flows removed
sp-2/0/0   ss1           0                  0

```

clear services nat mappings eim

| | |
|---------------------------------|---|
| Syntax | clear services nat mappings eim
<b4address <i>b4address/prefix</i> >
<subscriber <i>subscriber-ipv4-address</i> > |
| Release Information | Command introduced in Junos OS Release 14.1. |
| Description | Clear endpoint independent (EIM) and port control protocol (PCP) mappings . |
| Options | <p>none—Clear all EIM and PCP mappings.</p> <p>b4address <i>b4address/prefix</i>—(Optional) Clear EIM and PCP mappings for a particular subscriber <i>b4address/prefix</i></p> <p>internal-host <i>ipv4address/prefix</i>—(Optional) Clear EIM and PCP mappings matching the specified <i>b4address</i> and <i>internal-host</i>..</p> <p>port <i>port</i>—(Optional) Clear EIM and PCP mappings matching the specified <i>b4address</i>, <i>internal host</i>, and <i>port</i>.</p> <p>service-set <i>service-set</i>—(Optional) Clear EIM and PCP mappings for the specified <i>service set</i>.</p> <p>subscriber <i>subscriber-ipv4-address/prefix</i>—(Optional) Clear EIM and PCP mappings for a particular subscriber <i>ipv4-address/prefix</i></p> <ul style="list-style-type: none"> port <i>port</i>—(Optional) Clear EIM and PCP mappings matching the specified <i>ipv4-address/prefix</i> and <i>port</i>. service-set <i>service-set</i>—(Optional) Clear EIM and PCP mappings for the specified <i>service set</i>. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> show services nat mappings on page 2009 |
| List of Sample Output | clear services nat mappings eim on page 1843 |
| Output Fields | Table 57 on page 1842 lists the output fields for the clear services nat mappings eim command. Output fields are listed in the approximate order in which they appear. |

Table 57: clear services nat mappings eim Output Fields

| Field Name | Field Description |
|------------------|---|
| Interface | Name of a services interface. |
| Service set | Name of the service set from which flows are being cleared. |
| Mappings removed | Number of mappings removed. |

Table 57: clear services nat mappings eim Output Fields (*continued*)

| Field Name | Field Description |
|---------------|--------------------------|
| Flows removed | Number of flows removed. |

Sample Output

clear services nat mappings eim

```
user@host> clear services nat mappings eim
Interface  Service set  Mappings removed  Flows removed
sp-2/0/0   ss1          0                  0
```

clear services nat mappings pcg

| | |
|---------------------------------|---|
| Syntax | clear services nat mappings pcg
<b4address <i>b4address/prefix</i> >
<subscriber <i>subscriber-ipv4-address</i> > |
| Release Information | Command introduced in Junos OS Release 14.1. |
| Description | Clear NAT mappings for Port Control Protocol (PCP). |
| Options | <p>none—Clear all NAT PCP mappings.</p> <p>b4address <i>b4address/prefix</i>—(Optional) Clear NAT PCP mappings for a particular subscriber <i>b4address/prefix</i></p> <p>port <i>port</i>—(Optional) Clear NAT PCP mappings matching the specified <i>b4address</i> internal host, and port.</p> <p>service-set <i>service-set</i>—(Optional) Clear NAT PCP mappings for the specified service set.</p> <p>subscriber <i>ipv4-address/prefix</i>—(Optional) Clear NAT PCP mappings for a particular subscriber <i>ipv4-address/prefix</i></p> <p>port <i>port</i>—(Optional) Clear NAT PCP mappings matching the specified <i>ipv4-address/prefix</i>, and port.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> show services nat mappings on page 2009 |
| List of Sample Output | clear services nat mappings pcg on page 1845 |
| Output Fields | Table 58 on page 1844 lists the output fields for the clear services nat mappings pcg command. Output fields are listed in the approximate order in which they appear. |

Table 58: clear services nat mappings pcg Output Fields

| Field Name | Field Description |
|------------------|---|
| Interface | Name of a services interface. |
| Service set | Name of the service set from which flows are being cleared. |
| Mappings removed | Number of mappings removed. |
| Flows removed | Number of flows removed. |

Sample Output

clear services nat mappings pcg

```
user@host> clear services nat mappings pcg
Interface  Service set  Mappings removed  Flows removed
sp-2/0/0   ss1          0                 0
```

clear security pki ca-certificate

| | |
|---------------------------------|---|
| Syntax | clear security pki ca-certificate (all ca-profile <i>ca-profile-name</i>) |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | Delete certificate authority (CA) digital certificates from the router. |
| Options | all —Delete all CA digital certificates from the router.
ca-profile <i>ca-profile-name</i> —Delete the specified CA profile. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• request security pki ca-certificate enroll on page 1864• request security pki ca-certificate load on page 1865• show security pki ca-certificate on page 1930 |
| List of Sample Output | clear security pki ca-certificate all on page 1846 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear security pki ca-certificate all

```
user@host> clear security pki ca-certificate all
```


clear security pki certificate-request

| | |
|---------------------------------|--|
| Syntax | clear security pki certificate-request (all certificate-id <i>certificate-id-name</i>) |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | Delete manually generated local digital certificate requests from the router. |
| Options | <p>all—Delete all local digital certificate requests from the router.</p> <p>certificate-id <i>certificate-id-name</i>—Delete the specified local digital certificate and corresponding public/private key pair.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show security pki certificate-request on page 1934 |
| List of Sample Output | clear security pki certificate-request all on page 1847 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear security pki certificate-request all

```
user@host> clear security pki certificate-request all
```

clear security pki crt

| | |
|---------------------------------|---|
| Syntax | clear security pki crt (all ca-profile <i>ca-profile-name</i>) |
| Release Information | Command introduced in Junos 8.1 |
| Description | Delete certificate revocation lists (CRLs) from the router. |
| Options | all —Delete all CRLs from the router.

ca-profile <i>ca-profile-name</i> —Delete CRLs associated with the specified CA profile. |
| Required Privilege Level | clear |
| List of Sample Output | clear security pki crt ca-profile all on page 1848 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear security pki crt ca-profile all

```
user@host> clear security pki crt ca-profile all
```

clear security pki key-pair

| | |
|---------------------------------|--|
| Syntax | clear security pki key-pair (all certificate-id <i>certificate-id-name</i>) |
| Release Information | Command introduced in Junos OS Release 8.5. |
| Description | Clear public key infrastructure (PKI) key pair information for local digital certificates from the router. |
| Options | <p>all—Delete all local digital certificates, certificate requests, and the corresponding public and private key pairs from the router.</p> <p>certificate-id <i>certificate-id-name</i>—Delete the specified local digital certificate and corresponding public/private key pair.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• request security pki local-certificate enroll on page 1871• show security pki local-certificate on page 1938 |
| Output Fields | This command produces no output. |

Sample Output

```
user@host> clear security pki key pair
```

clear security pki local-certificate

| | |
|---------------------------------|--|
| Syntax | clear security pki local-certificate
<all certificate-id <i>certificate-id-name</i> system-generated> |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | Delete local digital certificates, certificate requests, and the corresponding public/private key pairs from the router. |
| Options | <p>all—(Optional) Delete all local digital certificates, certificate requests, and the corresponding public and private key pairs from the router.</p> <p>certificate-id <i>certificate-id-name</i>—(Optional) Delete the specified local digital certificate and corresponding public and private key pair.</p> <p>system-generated—(Optional) Auto-generated self-signed certificate.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• request security pki local-certificate enroll on page 1871• show security pki local-certificate on page 1938 |
| List of Sample Output | clear security pki local-certificate all on page 1850 |
| Output Fields | This command produces no output. |

Sample Output

clear security pki local-certificate all

```
user@host> clear security pki local-certificate all
```

clear services service-sets statistics integrity-drops

| | |
|-------------------------------------|--|
| Syntax | clear services service-sets statistics integrity-drops
<interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 13.3 |
| Description | Clear integrity-drops statistics for one adaptive services interface, for all adaptive services interfaces, or for one service-set. |
| Options | <p>none—Clear integrity-drops statistics for all configured adaptive service interfaces/
service-set.</p> <p>Service-set <i>service-set-name</i> —(Optional) Clear integrity-drops statistics for the specified
service-set</p> <p>interface <i>interface-name</i>—(Optional) Clear integrity-drops statistics for the specified
adaptive services interface.</p> |
| Required Privilege
Level | network |
| Related
Documentation | <ul style="list-style-type: none">• show services service-sets statistics packet-drops on page 2026• |

clear services service-sets statistics packet-drops

| | |
|---------------------------------|---|
| Syntax | clear services service-sets statistics packet-drops
<interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 7.4. |
| Description | Clear dropped-packet statistics for one adaptive services interface or for all adaptive services interfaces. |
| Options | none —Clear dropped-packet statistics for all configured adaptive services interfaces.

interface <i>interface-name</i> —(Optional) Clear dropped-packet statistics for the specified adaptive services interface. On M Series and T Series routers, the <i>interface-name</i> can be <i>ms-fpc/pic/port</i> , <i>sp-fpc/pic/port</i> or <i>rspnumber</i> . |
| Required Privilege Level | network |
| Related Documentation | <ul style="list-style-type: none">• show services service-sets statistics packet-drops on page 2026 |
| List of Sample Output | clear services service-sets statistics packet-drops on page 1852 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services service-sets statistics packet-drops

```
user@host> clear services service-sets statistics packet-drops interface sp-5/0/0
Flow collector interface: cp-5/0/0
Interface state: Collecting flows
Statistics cleared successfully
```

clear services service-sets statistics syslog

| | |
|---------------------------------|--|
| Syntax | clear services service-sets statistics syslog
<service-set <i>service-set-name</i> >
<interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 11.1. |
| Description | Clear system log statistics for one services interface or for all services interfaces, and for one named service set or all service sets on the interface or interfaces. |
| Options | <p>none—Clear system log for all configured services interfaces and their service sets.</p> <p>interface <i>interface-name</i>—(Optional) Clear system log statistics for the specified services interface. On M Series, MX Series, and T Series routers, the <i>interface-name</i> can be <i>ms-fpc/pic/port</i>, <i>sp-fpc/pic/port</i>, or <i>rspnumber</i>.</p> <p>service-set <i>service-set-name</i>—(Optional) Clear system log statistics for the specified services interface.</p> |
| Required Privilege Level | network |
| Related Documentation | <ul style="list-style-type: none"> • show services service-sets statistics syslog on page 2028 |
| List of Sample Output | clear services service-sets statistics syslog on page 1853 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services service-sets statistics syslog

```
user@host> clear services service-sets statistics syslog interface sp-5/0/0
Flow collector interface: cp-5/0/0
Interface state: Collecting flows
Statistics cleared successfully
```

clear services stateful-firewall flows

Syntax clear services stateful-firewall flows
<application-protocol *protocol*>
<destination-port *destination-port*>
<destination-prefix *destination-prefix*>
<interface *interface-name*>
<protocol *protocol*>
<service-set *service-set*>
<source-port *source-port*>
<source-prefix *source-prefix*>

Release Information Command introduced before Junos OS Release 7.4.

Description Clear stateful firewall flows. Issue this command to clear the stateful firewall flows for the specified option. The default option is "none", that is, to close all stateful firewall flows unless another option is specified.

Starting in Junos Release 14.1, the method for closing flows has changed. With the change, even for peak flows, the command prompt now returns to an active state after 30 seconds and the clear command completes in 90 to 120 seconds. In previous releases, closing peak flows could take as long as 4 minutes, after which the command prompt would return. Note too that during the first 30 seconds of issuing the command, the flows to be deleted remain visible in the **show services stateful-firewall flows** command output.

Options **none**—Clear all stateful firewall flows.

destination-port *destination-port*—(Optional) Clear stateful firewall flows for a particular destination port. The range of values is 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Clear stateful firewall flows for a particular destination prefix.

interface *interface-name*—(Optional) Clear stateful firewall flows for a particular interface. On M Series and T Series routers, the *interface-name* can be *ms-fpc/pic/port* or *rspnumber*.

protocol—(Optional) Clear stateful firewall flows for one of the following IP types:

- **number**—Numeric protocol value from 0 to 255.
- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-over-IP Encapsulation Protocol

- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Clear stateful firewall flows for a particular service set.

source-port *source-port*—(Optional) Clear stateful firewall flows for a particular source port. The range of values is from 0 through 65535.

source-prefix *source-prefix*—(Optional) Clear stateful firewall flows for a particular source prefix.

Required Privilege Level view

Related Documentation • [show services stateful-firewall flows on page 2052](#)

List of Sample Output [clear services stateful-firewall flows on page 1855](#)

Output Fields [Table 59 on page 1855](#) lists the output fields for the **clear services stateful-firewall flows** command. Output fields are listed in the approximate order in which they appear.

Table 59: clear services stateful-firewall flows Output Fields

| Field Name | Field Description |
|---------------------|---|
| Interface | Name of an adaptive services interface. |
| Service set | Name of the service set from which flows are being cleared. |
| Conv removed | Number of conversations removed. |

Sample Output

clear services stateful-firewall flows

```

user@host> clear services stateful-firewall flows
Interface  Service set                               Conv removed
ms-0/3/0   svc_set_trust                             0
ms-0/3/0   svc_set_untrust                           0

```

clear services stateful-firewall sip-call

Syntax clear services stateful-firewall sip-call
<application-protocol *protocol*>
<destination-port *destination-port*>
<destination-prefix *destination-prefix*>
<interface *interface-name*>
<protocol *protocol*>
<service-set *service-set*>
<source-port *source-port*>
<source-prefix *source-prefix*>

Release Information Command introduced in Junos OS Release 7.4.

Description Clear Session Initiation Protocol (SIP) call information in stateful firewall flows.

Options **none**—Clear stateful firewall statistics for all interfaces and all service sets.

application-protocol—(Optional) Clear information about one of the following application protocols:

- **bootp**—(SIP only) Bootstrap protocol
- **dce-rpc**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—(SIP only) Domain Name System protocol
- **exec**—(SIP only) Exec
- **ftp**—(SIP only) File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol
- **iiop**—Internet Inter-ORB Protocol
- **login**—Login
- **netbios**—NetBIOS
- **netshow**—NetShow
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol
- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **shell**—Shell
- **sip**—Session Initiation Protocol

- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

destination-port *destination-port*—(Optional) Clear information for a particular destination port. The range of values is 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Clear information for a particular destination prefix.

interface *interface-name*—(Optional) Clear information for a particular adaptive services interface. On M Series and T Series routers, the *interface-name* can be **sp-fpc/pic/port** or **rspnumber**.

protocol—(Optional) Clear information about one of the following IP types:

- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol
- **ipv6**—IPv6 within IP
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Clear information for a particular service set.

source-port *source-port*—(Optional) Clear information for a particular source port. The range of values is 0 to 65535.

source-prefix *source-prefix*—(Optional) Clear information for a particular source prefix.

Required Privilege Level view

**Related
Documentation**

- [show services stateful-firewall sip-call on page 2058](#)

List of Sample Output

[clear services stateful-firewall sip-call on page 1858](#)

Output Fields

[Table 60 on page 1858](#) lists the output fields for the **clear services stateful-firewall sip-call** command. Output fields are listed in the approximate order in which they appear.

Table 60: clear services stateful-firewall sip-call Output Fields

| Field Name | Field Description |
|-------------------|---|
| Interface | Name of an adaptive services interface. |
| Service set | Name of the service set from which flows are being cleared. |
| SIP calls removed | Number of SIP calls removed. |

Sample Output**clear services stateful-firewall sip-call**

```
user@host> clear services stateful-firewall sip-call
Interface  Service set      SIP calls removed
sp-0/3/0   test_sip_777     1
```

clear services stateful-firewall sip-register

Syntax clear services stateful-firewall sip-register
 <application-protocol *protocol*>
 <destination-port *destination-port*>
 <destination-prefix *destination-prefix*>
 <interface *interface-name*>
 <protocol *protocol*>
 <service-set *service-set*>
 <source-port *source-port*>
 <source-prefix *source-prefix*>

Release Information Command introduced in Junos OS Release 7.4.

Description Clear Session Initiation Protocol (SIP) register information in stateful firewall flows.

Options **application-protocol**—(Optional) Clear information about one of the following application protocols:

- **bootp**—(SIP only) Bootstrap protocol
- **dce-rpc**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—(SIP only) Domain Name System protocol
- **exec**—(SIP only) Exec
- **ftp**—(SIP only) File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol
- **iiop**—Internet Inter-ORB Protocol
- **login**—Login
- **netbios**—NetBIOS
- **netshow**—NetShow
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol
- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **shell**—Shell
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet

- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

destination-port *destination-port*—(Optional) Clear information for a particular destination port. The range of values is 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Clear information for a particular destination prefix.

interface *interface*—(Optional) Clear information about a particular interface. On M Series and T Series routers, the *interface-name* can be **sp-fpc/pic/port** or **rspnumber**.

protocol—(Optional) Clear information about one of the following IP types:

- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol
- **ipv6**—IPv6 within IP
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Clear information for a particular service set.

source-port *source-port*—(Optional) Clear information for a particular source port. The range of values is 0 through 65535.

source-prefix *source-prefix*—(Optional) Clear information for a particular source prefix.

Required Privilege
Level

view

Related
Documentation

- [show services stateful-firewall sip-register on page 2063](#)

List of Sample Output

[clear services stateful-firewall sip-register on page 1861](#)

Output Fields Table 61 on page 1861 lists the output fields for the **clear services stateful-firewall sip-register** command. Output fields are listed in the approximate order in which they appear.

Table 61: clear services stateful-firewall sip-register Output Fields

| Field Name | Field Description |
|---------------------------------|---|
| Interface | Name of an adaptive services interface. |
| Service set | Name of the service set from which flows are being cleared. |
| SIP registration removed | Number of SIP registers removed. |

Sample Output

clear services stateful-firewall sip-register

```
user@host> clear services stateful-firewall sip-register
Interface  Service set      SIP registration removed
sp-0/3/0   test_sip_777     1
```

clear services stateful-firewall statistics


| | |
|--------------------------|--|
| Syntax | clear services stateful-firewall statistics
<interface <i>interface-name</i> >
<service-set <i>service-set</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear stateful firewall statistics. |
| Options | <p>none—Clear stateful firewall statistics for all interfaces and all service sets.</p> <p>interface <i>interface-name</i>—(Optional) Clear stateful firewall statistics for the specified interface. On M Series and T Series routers, the <i>interface-name</i> can be <i>ms-fpc/pic/port</i> or <i>rspnumber</i>.</p> <p>service-set <i>service-set</i>—(Optional) Clear stateful firewall statistics for the specified service set.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show services stateful-firewall statistics on page 2067 |
| List of Sample Output | clear services stateful-firewall statistics on page 1862 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services stateful-firewall statistics

```
user@host> clear services stateful-firewall statistics
```


request interface (revert | switchover) (Adaptive Services)

| | |
|--|--|
| Syntax | request interface (revert switchover) (<i>rspnumber</i> <i>rlsqnumber</i>) |
| Release Information | Command introduced before Junos OS Release 7.4.
Support for rlsq interfaces added in Junos OS Release 7.6. |
| Description | (M Series and T Series routers only) Manually revert to the primary adaptive services interface or link services IQ interface, or to switch from the primary to the secondary interface. |
| <div>  NOTE: All rlsq switchover or revert operations are allowed from the rlsqnumber level only and not for individual channelized interfaces (rlsqnumber:unit). </div> | |
| <p>On an aggregated Ethernet interface with link protection enabled, use the request interface (revert switchover) (Aggregated Ethernet Link Protection) operational command to manually revert egress traffic from the designated backup link to the designated primary link, or to manually switch egress traffic from the primary link to the backup link. For information about this command, see <i>request interface (revert switchover) (Aggregated Ethernet Link Protection)</i>.</p> | |
| Options | <p>(revert switchover)—The revert keyword restores active processing to the primary adaptive services (sp) or link services IQ (lsq) interface. The switchover keyword transfers active processing to the secondary (backup) interface.</p> <p>rspnumber—Redundant adaptive services interface name.</p> <p>rlsqnumber—Redundant link services IQ interface name.</p> |
| Required Privilege Level | view |
| List of Sample Output | request interface revert on page 1863
request interface switchover on page 1863 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request interface revert

```
user@host> request interface revert rlsq0
request succeeded
```

request interface switchover

```
user@host> request interface switchover rlsq0
error: rlsq0: already on secondary
```

request security pki ca-certificate enroll

| | |
|---------------------------------|--|
| Syntax | request security pki ca-certificate enroll ca-profile <i>ca-profile-name</i> |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | Request a digital certificate from a certificate authority (CA) online by using the Simple Certificate Enrollment Protocol (SCEP). |
| Options | ca-profile <i>ca-profile-name</i> —CA profile name. |
| Required Privilege Level | maintenance |
| Related Documentation | <ul style="list-style-type: none">• clear security pki ca-certificate on page 1846• show security pki ca-certificate on page 1930 |
| List of Sample Output | request security pki ca-certificate enroll on page 1864 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request security pki ca-certificate enroll

```
user@host> request security pki ca-certificate enroll ca-profile entrust
Received following certificates:
Certificate: C=us, O=juniper, CN=First Officer
Fingerprint: 46:71:15:34:f0:a6:41:76:65:81:33:4f:68:47:c4:df:78:b8:e3:3f
Certificate: C=us, O=juniper, CN=First Officer
Fingerprint: bc:78:87:9b:a7:91:13:20:71:db:ac:b5:56:71:42:ad:1a:b6:46:17
Certificate: C=us, O=juniper
Fingerprint: 00:8e:6f:58:dd:68:bf:25:0a:e3:f9:17:70:d6:61:f3:53:a7:79:10
Do you want to load the above CA certificate ? [yes,no] (no) yes
```

request security pki ca-certificate load

| | |
|---------------------------------|---|
| Syntax | <code>request security pki ca-certificate load ca-profile <i>ca-profile-name</i> filename <i>path/filename</i></code> |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | Manually load a certificate authority (CA) digital certificate from a specified location. |
| Options | <p>ca-profile <i>ca-profile-name</i>—Load the specified CA profile.</p> <p>filename <i>path/filename</i>—Directory location and filename of the CA digital certificate.</p> |
| Required Privilege Level | maintenance |
| Related Documentation | <ul style="list-style-type: none"> • clear security pki ca-certificate on page 1846 • show security pki ca-certificate on page 1930 |
| List of Sample Output | request security pki ca-certificate load on page 1865 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request security pki ca-certificate load

```
user@host> request security pki ca-certificate load ca-profile ca-private filename pki-file
```

request security pki ca-certificate verify

| | |
|---------------------------------|---|
| Syntax | <code>request security pki ca-certificate verify ca-profile <i>ca-profile-name</i></code> |
| Release Information | Command introduced in Junos OS Release 8.5. |
| Description | Verify the digital certificate installed for the specified certificate authority (CA). |
| Options | ca-profile <i>ca-profile-name</i> —Name of the local digital certificate identifier. |
| Required Privilege Level | maintenance |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

You receive the following response before the certificate revocation list (CRL) is downloaded:

```
request security pki ca-certificate verify ca-profile ca1 (CRL not downloaded)
user@host> request security pki ca-certificate verify ca-profile ca1
```

```
CA certificate ca1: CRL verification in progress. Please check the PKId debug
logs for completion status
```

request security pki crt load

| | |
|---------------------------------|---|
| Syntax | <code>request security pki crt load ca-profile <i>ca-profile-name</i> filename <i>path/filename</i></code> |
| Release Information | Command introduced in Junos OS Release 8.1. |
| Description | Manually install a certificate revocation list (CRL) on the router from a specified location. |
| Options | <code>ca-profile <i>ca-profile-name</i></code> —Load the specified certificate authority (CA) profile.
<code>filename <i>path/filename</i></code> —Directory location and filename of the CRL. |
| Required Privilege Level | maintenance |
| List of Sample Output | request security pki crt load on page 1867 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request security pki crt load

```
user@host> request security pki crt load ca-profile ca-private filename pki-file
```

request security pki generate-certificate-request

| | |
|---------------------------------|--|
| Syntax | <code>request security pki generate-certificate-request certificate-id <i>certificate-id-name</i>
domain-name <i>domain-name</i> subject <i>subject-distinguished-name</i>
<email <i>email-address</i>>
<filename (<i>path</i> <i>terminal</i>)>
<ip-address <i>ip-address</i>></code> |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | Manually generate a local digital certificate request in the Public-Key Cryptography Standards #10 (PKCS-10) format. |
| Options | <p>certificate-id <i>certificate-id-name</i>—Name of the local digital certificate and the public/private key pair.</p> <p>domain-name <i>domain-name</i>—Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.</p> <p>subject <i>subject-distinguished-name</i>—Distinguished name format that contains the common name, department, company name, state, and country:</p> <ul style="list-style-type: none">• CN—Common name• OU—Organizational unit name• O—Organization name• ST—State• C—Country <p>email <i>email-address</i>—(Optional) E-mail address of the certificate holder.</p> <p>filename (<i>path</i> <i>terminal</i>)—(Optional) Location where the local digital certificate request should be placed or the login terminal.</p> <p>ip-address <i>ip-address</i>—(Optional) IP address of the router.</p> |
| Required Privilege Level | maintenance |
| Related Documentation | <ul style="list-style-type: none">• clear security pki certificate-request on page 1847• show security pki certificate-request on page 1934 |
| List of Sample Output | request security pki generate-certificate-request on page 1869 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request security pki generate-certificate-request

```
user@host> request security pki generate-certificate-request certificate-id local-entrust2
domain-name router2.juniper.net filename entrust-req2 subject cn=router2.juniper.net
```

Generated certificate request

-----BEGIN CERTIFICATE REQUEST-----

```
MIIBOTCCAQoCAQAwGjEYMBYGA1UEAxMPdHxLmp1bm1wZXIubmV0MIGfMA0GCSqG
SIb3DQEBAQUAA4GNADCBiQKBgQCiUFk1Qws1Ud+AqN5DDxRs2kVyKEhh9qoVFnz+
Hz4c9vsy3B8E1wTJ1kmIt2cB3yi fB6zePd+6WYpf57Crwre7YqPkiXM31F6z3YjX
H+1BPNbCxNWYvyrnSyVYDbFj8o0Xyqog8ACDFVL2JBWrPNBYy7imq/K9soDBbAs6
5hZqqwIDAQABoEcwRQYJKoZIhvcNAQkOMTgwNjA0BGNVHQ8BAf8EBAMCB4AwJAYD
VR0RAQH/BBowGIIWdHxLmVuZ2xhYi5qdW5pcGVyLm5ldDANBgkqhkiG9w0BAQQF
AAOBgQBc2rq1v5S0QXH7LCb/FdqAL8ZM6GoaN5d6cGwq4bB6a7UQFgtoH406gQ3G
3iH0Zfz4xMIBpJYuGd1dkqgvcDoH3AgTsLkfn7Wi3x5H2qeQVs9bvL4P5nvEZLND
EIMUHwteo1ZCiZ70f09Fer9cXWHSQs1UtXtgPqQJy2xIeImLgw==
```

-----END CERTIFICATE REQUEST-----

Fingerprint:

```
0d:90:b8:d2:56:74:fc:84:59:62:b9:78:71:9c:e4:9c:54:ba:16:97 (sha1)
1b:08:d4:f7:90:f1:c4:39:08:c9:de:76:00:86:62:b8 (md5)
```

request security pki generate-key-pair

| | |
|---------------------------------|---|
| Syntax | request security pki generate-key-pair certificate-id <i>certificate-id-name</i>
<size (512 1024 2048)> |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | Generate a Public Key Infrastructure (PKI) public and private key pair for a local digital certificate. |
| Options | certificate-id <i>certificate-id-name</i> —Name of the local digital certificate and the public/private key pair.

size —(Optional) Key pair size. The key pair size can be 512 , 1024 , or 2048 bits. |
| Required Privilege Level | maintenance |
| List of Sample Output | request security pki generate-key-pair on page 1870 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request security pki generate-key-pair

```
user@host> request security pki generate-key-pair certificate-id billy size 2048
Generated key pair billy, key size 2048 bits
```


request security pki local-certificate enroll

| | |
|---------------------------------|--|
| Syntax | request security pki local-certificate enroll <i>ca-profile ca-profile-name</i>
<i>certificate-id certificate-id-name</i> challenge-password <i>password</i> domain-name
<i>domain-name</i> subject <i>subject-distinguished-name</i>
<email <i>email-address</i> >
<ip-address <i>ip-address</i> > |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | Request that a certificate authority (CA) enroll and install a local digital certificate online by using the Simple Certificate Enrollment Protocol (SCEP). |
| Options | <p>ca-profile <i>ca-profile-name</i>—CA profile name.</p> <p>certificate-id <i>certificate-id-name</i>—Name of the local digital certificate and the public/private key pair.</p> <p>challenge-password <i>password</i>—Password set by the administrator and normally obtained from the SCEP enrollment webpage of the CA. The password is 16 characters in length.</p> <p>domain-name <i>domain-name</i>—Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.</p> <p>subject <i>subject-distinguished-name</i>—Distinguished name format that contains the common name, department, company name, state, and country:</p> <ul style="list-style-type: none"> • CN—Common name • OU—Organizational unit name • O—Organization name • ST—State • C—Country <p>email <i>email-address</i>—(Optional) E-mail address of the certificate holder.</p> <p>ip-address <i>ip-address</i>—(Optional) IP address of the router.</p> |
| Required Privilege Level | maintenance |
| Related Documentation | <ul style="list-style-type: none"> • show security pki local-certificate on page 1938 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
user@host> request security pki local-certificate enroll certificate-id r3-entrust-scep ca-profile  
entrust domain-name router3.juniper.net subject "CN=router3,OU=Engineering,O=juniper,C=US"  
challenge-password 123
```

Certificate enrollment has started. To view the status of your enrollment, check the public key infrastructure log (pkid) log file at /var/log/pkid. Please save the challenge-password for revoking this certificate in future. Note that this password is not stored on the router.

request security pki local-certificate generate-self-signed

| | |
|---------------------------------|---|
| Syntax | <code>request security pki local-certificate generate-self-signed certificate-id <i>certificate-id-name</i>
domain-name <i>domain-name</i> ip-address <i>ip-address</i> email <i>email-address</i>
subject <i>subject-distinguished-name</i></code> |
| Release Information | Command introduced in Junos OS Release 9.1. |
| Description | Manually generate a self-signed certificate for the given distinguished name. |
| Options | <p>certificate-id <i>certificate-id-name</i>—Name of the local digital certificate and the public/private key pair.</p> <p>domain-name <i>domain-name</i>—Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.</p> <p>email <i>email-address</i>—E-mail address of the certificate holder.</p> <p>ip-address <i>ip-address</i>—IP address of the router.</p> <p>subject <i>subject-distinguished-name</i>—Distinguished name format that contains the common name, department, company name, state, and country:</p> <ul style="list-style-type: none"> • CN—Common name • OU—Organizational unit name • O—Organization name • ST—State • C—Country |
| Required Privilege Level | maintenance
security |
| Related Documentation | <ul style="list-style-type: none"> • show security pki local-certificate on page 1938 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
user@host> request security pki local-certificate generate-self-signed certificate-id self-cert
subject cn=abc domain-name juniper.net email mholmes@juniper.net
Self-signed certificate generated and loaded successfully
```

request security pki local-certificate load

| | |
|---------------------------------|---|
| Syntax | <code>request security pki local-certificate load certificate-id <i>certificate-id-name</i> filename <i>path</i></code> |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | Manually load a local digital certificate from a specified location. |
| Options | <p>certificate-id <i>certificate-id-name</i>—Name of the public/private key pair mapped to the local digital certificate.</p> <p>filename <i>path/filename</i>—Directory location and filename of the local digital certificate provided by the CA.</p> |
| Required Privilege Level | maintenance |
| List of Sample Output | request security pki local-certificate load on page 1874 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request security pki local-certificate load

```
user@host> request security pki local-certificate load filename /tmp/router2-cert certificate-id
local-entrust2
Local certificate local-entrust2 loaded successfully
```

request security pki local-certificate verify

| | |
|---------------------------------|--|
| Syntax | <code>request security pki local-certificate verify certificate-id <i>certificate-id-name</i></code> |
| Release Information | Command introduced in Junos OS Release 8.5. |
| Description | Verify the validity of the local digital certificate identifier. |
| Options | <code>certificate-id <i>certificate-id-name</i></code> —Display the specified certificate identifier name. |
| Required Privilege Level | maintenance |
| Related Documentation | <ul style="list-style-type: none"> • show security pki local-certificate on page 1938 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

You receive the following response before the certificate revocation list (CRL) is downloaded:

```
request security pki local-certificate verify certificate-id bme1 (not downloaded)
user@host> request security pki local-certificate verify certificate-id bme1
```

```
Local certificate bme1: CRL verification in progress. Please check the PKId debug
logs for completion status
```

You receive the following response after the certificate revocation list (CRL) is downloaded:

```
request security pki local-certificate verify certificate bme1 (downloaded)
user@host> request security pki local-certificate verify certificate-id bme1
Local certificate bme1 verification success
```

request services ipsec-vpn ipsec switch tunnel

| | |
|--------------------------|--|
| Syntax | request services ipsec-vpn ipsec switch tunnel local-gateway <i>address</i> remote-gateway <i>address</i>
<routing-instance <i>instance-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4.
routing-instance option added in Release 8.1. |
| Description | (Adaptive services interface only) Manually switch between primary and backup IP Security (IPsec) tunnels. |
| Options | local-gateway <i>address</i> —Gateway address of the local system.

remote-gateway <i>address</i> —Gateway address of the remote system.

routing-instance <i>instance-name</i> —(Optional) VRF instance associated with local gateway address. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show services ipsec-vpn ipsec security-associations on page 1965 |
| List of Sample Output | request services ipsec-vpn ipsec switch tunnel on page 1876 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request services ipsec-vpn ipsec switch tunnel

```
user@host> request services ipsec-vpn ipsec switch tunnel local-gateway 10.1.1.1 remote gateway 10.100.10.1
```

show interfaces (Adaptive Services)

| | |
|---------------------------------|--|
| Syntax | <pre>show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display status information about the specified adaptive services interface. |
| Options | <p><i>interface-type</i>—On M Series and T Series routers, the interface type is sp-<i>fpc/pic/port</i>.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show interfaces (Adaptive Services) on page 1882</p> <p>show interfaces brief (Adaptive Services) on page 1882</p> <p>show interfaces detail (Adaptive Services) on page 1882</p> <p>show interfaces extensive (Adaptive Services) on page 1883</p> |
| Output Fields | Table 62 on page 1877 lists the output fields for the show interfaces (adaptive services and redundant adaptive services) command. Output fields are listed in the approximate order in which they appear. |

Table 62: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|--|------------------------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface's index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |

Table 62: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (continued)

| Field Name | Field Description | Level of Output |
|--------------------------------|--|------------------------------|
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Type | Encapsulation being used on the interface. | All levels |
| Link-level type | Encapsulation being used on the physical interface. | All levels |
| MTU | MTU size on the physical interface. | All levels |
| Clocking | Reference clock source: can be Internal or External . | All levels |
| Speed | Speed at which the interface is running. | All levels |
| Device flags | Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Link type | Physical interface link type: Full-Duplex or Half-Duplex . | detail extensive none |
| Link flags | Information about the link. Possible values are described in the "Link Flags" section under <i>Common Output Fields Description</i> . | detail extensive none |
| Physical info | Information about the physical interface. | detail extensive |
| Hold-times | Current interface hold-time up and hold-time down, in milliseconds. | detail extensive |
| Current address | Configured MAC address. | detail extensive none |
| Hardware address | MAC address of the hardware. | detail extensive none |
| Alternate link address | Backup address of the link. | detail extensive none |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) . | detail extensive none |
| Input Rate | Input rate in bits per second (bps) and packets per second (pps). | None specified |
| Output Rate | Output rate in bps and pps. | None specified |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |

Table 62: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (continued)

| Field Name | Field Description | Level of Output |
|---------------------------|---|------------------------------|
| Traffic statistics | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. | detail extensive |
| Input errors | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Frames received smaller than the runt threshold. • Giants—Frames received larger than the giant threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • Resource errors—Sum of transmit drops. | extensive |
| Output errors | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • MTU errors—Number of packets larger than the MTU threshold. • Resource errors—Sum of transmit drops. | extensive |
| Logical Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |

Table 62: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------------------------|---|------------------------------|
| Flags | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Encapsulation | Encapsulation on the logical interface. | All levels |
| Input packets | Number of packets received on the logical interface. | None specified |
| Output packets | Number of packets transmitted on the logical interface. | None specified |
| Traffic statistics | <p>Number and rate of bytes and packets received and transmitted on the logical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. | detail extensive |
| Local statistics | Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Transit statistics | Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes generally less than 1 second for the counter to stabilize. | detail extensive |
| <i>protocol-family</i> | Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed. | brief |
| Protocol | Protocol family configured on the logical interface, such as iso , inet6 , mpls . | detail extensive none |
| MTU | MTU size on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route table | Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 . | detail extensive |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Addresses, Flags | Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Destination | IP address of the remote side of the connection. | detail extensive none |
| Local | IP address of the logical interface. | detail extensive none |

Table 62: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields
(continued)

| Field Name | Field Description | Level of Output |
|-------------------|---|------------------------------|
| Broadcast | Broadcast address. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |

Sample Output

show interfaces (Adaptive Services)

```
user@host> show interfaces sp-1/2/0
Physical interface: sp-1/2/0, Enabled, Physical link is Up
  Interface index: 147, SNMP ifIndex: 72
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : 2006-03-06 11:37:18 PST (00:57:29 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

Logical interface sp-1/2/0.16383 (Index 68) (SNMP ifIndex 73)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
  Input packets : 3057
  Output packets: 3044
  Protocol inet, MTU: 9192
  Flags: Receive-options, Receive-TTL-Exceeded
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.0.0.34, Local: 10.0.0.1
```

show interfaces brief (Adaptive Services)

```
user@host> show interfaces sp-1/2/0 brief
Physical interface: sp-1/2/0, Enabled, Physical link is Up
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000

Logical interface sp-1/2/0.16383
  Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
  inet 10.0.0.1      --> 10.0.0.34
```

show interfaces detail (Adaptive Services)

```
user@host> show interfaces sp-1/2/0 detail
Physical interface: sp-1/2/0, Enabled, Physical link is Up
  Interface index: 147, SNMP ifIndex: 72, Generation: 30
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link type      : Full-Duplex
  Link flags     : None
  Physical info   : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2006-03-06 11:37:18 PST (00:57:56 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :          125147          0 bps
    Output bytes :         1483113          0 bps
    Input packets:           3061          0 pps
    Output packets:          3048          0 pps
```

```

Logical interface sp-1/2/0.16383 (Index 68) (SNMP ifIndex 73) (Generation 7)
Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
Traffic statistics:
  Input bytes :          125147
  Output bytes :        1483113
  Input packets:          3061
  Output packets:        3048
Local statistics:
  Input bytes :          125147
  Output bytes :        1483113
  Input packets:          3061
  Output packets:        3048
Transit statistics:
  Input bytes :           0          0 bps
  Output bytes :           0          0 bps
  Input packets:           0          0 pps
  Output packets:          0          0 pps
Protocol inet, MTU: 9192, Generation: 20, Route table: 1
Flags: Receive-options, Receive-TTL-Exceeded
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.0.0.34, Local: 10.0.0.1, Broadcast: Unspecified,
  Generation: 22

```

show interfaces extensive (Adaptive Services)

```

user@host> show interfaces sp-1/2/0 extensive
Physical interface: sp-1/2/0, Enabled, Physical link is Up
Interface index: 147, SNMP ifIndex: 72, Generation: 30
Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
Clocking: Unspecified, Speed: 800mbps
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
Link type : Full-Duplex
Link flags : None
Physical info : Unspecified
Hold-times : Up 0 ms, Down 0 ms
Current address: Unspecified, Hardware address: Unspecified
Alternate link address: Unspecified
Last flapped : 2006-03-06 11:37:18 PST (00:58:40 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes :          125547          0 bps
  Output bytes :        1483353          0 bps
  Input packets:          3065          0 pps
  Output packets:        3052          0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 2, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0

Logical interface sp-1/2/0.16383 (Index 68) (SNMP ifIndex 73) (Generation 7)
Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
Traffic statistics:
  Input bytes :          125547
  Output bytes :        1483353
  Input packets:          3065
  Output packets:        3052
Local statistics:

```

```
Input bytes :          125547
Output bytes :         1483353
Input packets:           3065
Output packets:          3052
Transit statistics:
Input bytes :              0          0 bps
Output bytes :              0          0 bps
Input packets:              0          0 pps
Output packets:             0          0 pps
Protocol inet, MTU: 9192, Generation: 20, Route table: 1
Flags: Receive-options, Receive-TTL-Exceeded
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.0.0.34, Local: 10.0.0.1, Broadcast: Unspecified,
Generation: 22
```

show interfaces (Link Services IQ)

| | |
|---------------------------------|--|
| Syntax | <pre>show interfaces lsq-<i>fpc/pic/port</i> <brief detail extensive terse> <descriptions> <l2-statistics> <media> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>l2-statistics option introduced with Junos OS Release 12.1.</p> |
| Description | (M Series, MX Series, and T Series routers only) Display status information about the specified link services intelligent queuing (IQ) interface. |
| Options | <p>lsq-<i>fpc/pic/port</i>—Display standard status information about the specified link services IQ interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>l2-statistics—(Optional) Display Layer 2 queue statistics for Multilink Point-to-Point Protocol (MLPPP), FRF.15, and FRF.16 bundles.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Additional Information | Link services IQ interfaces are similar to link services interfaces. The important difference is that link services IQ interfaces fully support Junos OS class-of-service (CoS) components. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • Link and Multilink Services Overview on page 711 • Multilink Interfaces on Channelized MICs Overview on page 715 |
| List of Sample Output | <p>show interfaces extensive (MLPPP on Link Services IQ) on page 1900</p> <p>show interfaces extensive (Multiclass MLPPP on Link Services IQ) on page 1901</p> <p>show interfaces extensive (MLPPP on Link Services IQ Bundle) on page 1903</p> <p>show interfaces extensive (MFR on Link Services IQ Bundle) on page 1904</p> <p>show interfaces (Multiclass MLPPP on Link Services IQ) on page 1906</p> |

Output Fields Table 63 on page 1886 lists the output fields for the **show interfaces** (link services IQ) command. Output fields are listed in the approximate order in which they appear.

Table 63: show interfaces (Link Services IQ) Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|--|------------------------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| Link-level type | Encapsulation being used on the physical interface:
Multilink-Frame-Relay-UNI-NNI Multilink-Frame-Relay-UNI-NNI (default),
LinkService , Frame-relay , Frame-relay-ccc , or Frame-relay-tcc . | All levels |
| MTU | Maximum transmission unit size on the physical interface. | All levels |
| Device flags | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) . | detail extensive none |

Table 63: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|--|------------------------------|
| Multilink Frame Relay UNI NNI bundle options | <p>(Multilink Frame Relay UNI NNI only) Configured information about Multilink Frame Relay bundle options.</p> <ul style="list-style-type: none"> • Device type—DCE (data communication equipment) or DTE (data terminal equipment). • MRRU—Configured size of the maximum received reconstructed unit (MRRU): 1500 to 4500 bytes. The default is 1524 bytes. • Bandwidth—Speed at which the interface is running. • Fragmentation threshold—Configured fragmentation threshold: 128 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation. • Red differential delay limit—Red differential delay limit among bundle links has been reached, indicating an action will occur. • Yellow differential delay limit—Yellow differential delay among bundle links has been reached, indicating a warning will occur. • Red differential delay action—Type of actions taken when the red differential delay exceeds the red limit: <i>Disable link transmit</i> or <i>Remove link from service</i>. • Link layer overhead—Percentage of bundle bandwidth to be set aside for link layer overhead. • Reassembly drop timer—Drop timeout value to provide a recovery mechanism if individual links in the link services bundle drop one or more packets: 1 through 127 milliseconds. By default, the drop timeout parameter is 0 (disabled). A value under 5 ms is not recommended. • Links needed to sustain bundle—Minimum number of links to sustain the bundle: 1 through 8. • LIP Hello timer—Link Interleaving Protocol hello timer: 1 through 180 seconds. <ul style="list-style-type: none"> • Acknowledgement timer—Maximum period to wait for an add link acknowledgement, hello acknowledgement, or remove link acknowledgement: 1 through 10 seconds. • Acknowledgement retries—Number of retransmission attempts to be made for consecutive hello or remove link messages after the expiration of the acknowledgement timer: 1 through 5. | detail extensive none |

Table 63: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|---|-----------------------|
| Multilink Frame Relay UNI NNI bundle options (continued) | <ul style="list-style-type: none"> • Bundle class—Bundle class ID. • LMI type—Multilink Frame Relay UNI NNI LMI type: ANSI, Q.933 ANNEX A, or Consortium. <ul style="list-style-type: none"> • T391 LIV polling timer—Multilink Frame Relay UNI NNI Full status polling counter: 1 through 255, with a default value of 6. • T392 polling verification timer—Multilink Frame Relay UNI NNI LMI error threshold. The number of errors required to bring down the link, within the event count specified by <i>N393</i>. The range is 1 through 10, with a default value of 3. • N391 full status polling count—Multilink Frame Relay UNI NNI Full status polling counter: 1 through 255. • N392 error threshold—Multilink Frame Relay UNI NNI LMI error threshold: 1 through 10. • N393 monitored event count—Multilink Frame Relay UNI NNI LMI monitored event count: 1 through 10, with a default value of 4. • Consortium LMI Settings <ul style="list-style-type: none"> • n391dte—DTE full status polling interval in seconds: 1 through 255. • n392dce—DCE error threshold: 1 through 10. • n392dte—DTE error threshold: 1 through 10. • n393dce—DCE monitored event count: 1 through 10. • n393dte—DTE monitored event count: 1 through 10. • t391dte—DTE polling verification timer (in seconds): 5 through 30. • t392dce—DCE polling verification timer (in seconds): 5 through 30. | detail extensive none |
| LMI | <p>Local Management Interface packet statistics:</p> <ul style="list-style-type: none"> • Input—Number of packets arriving on the interface (nn) and timestamp of the most recent packet arrival, in the format:
 Input: nn (last seen hh:mm:ss ago) • Output—Number of packets sent out on the interface (nn) and how much time has passed since the last packet was sent, in the format:
 Output: nn (last seen hh:mm:ss ago) | detail extensive none |
| DTE Statistics | <p>Statistics about information transferred from the data terminal equipment (DTE) to the data communications equipment (DCE).</p> <ul style="list-style-type: none"> • Enquiries sent—Number of link status enquiries sent from the DTE to the DCE. • Full enquiries sent—Number of full enquiries sent from the DTE to the DCE. • Enquiry responses received—Number of enquiry responses received by the DCE from the DTE. • Full enquiry responses received—Number of full enquiry responses received by DCE from the DTE. | detail extensive none |

Table 63: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|--|------------------------------|
| DCE Statistics | <p>Statistics about information transferred from the DCE to the DTE.</p> <ul style="list-style-type: none"> • Enquiries received—Number of enquiries received by the DCE from the DTE. • Full enquiries received—Number of full enquiries received by the DCE from the DTE. • Enquiry responses sent—Number of enquiry responses sent from the DCE to the DTE. • Full enquiry responses sent—Number of full enquiry responses sent from the DCE to the DTE. | detail extensive none |
| Common Statistics | <p>Statistics about messages sent between the DTE and the DCE.</p> <ul style="list-style-type: none"> • Unknown messages received—Number of received packets that do not fall into any other category. • Asynchronous updates received—Number of link status peer changes received. • Out-of-sequence packets received—Number of packets for which the sequence of the packets received is different from the expected sequence. • Keepalive responses timed out—Number of keepalive responses that time out when no Local Management Interface (LMI) packet was reported for n392dte or n393dce intervals. (See <i>LMI settings</i>.) | |
| Traffic statistics | <p>Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the Packet Forwarding Engine (PFE). Input traffic refers to the fragments received by the ingress PFE, which get assembled into Layer 3 input packets. Output packets refer to the IP packets transmitted out of the ingress PFE to the LSQ, which get segmented into output fragments.</p> | detail extensive |
| DLCInn | <p>Data-link connection identifier (DLCI) number of the logical interface. The following information is displayed.</p> <ul style="list-style-type: none"> • Flags—Values are: <ul style="list-style-type: none"> • Active—Set when the link is active and the DTE and DCE are exchanging information. • Down—Set when the link is active, but no information is received from the DTE. • DCE unconfigured—Set when the corresponding DLCI in the DCE is not configured. • Configured—Set when the corresponding DLCCI is configured. • DCE-Configured—Displayed when the command is issued from the DTE. | |
| DLCI Statistics | <p>(Frame Relay) Data-link connection identifier (DLCI) statistics.</p> <ul style="list-style-type: none"> • Active DLCI—Number of active DLCIs. • Inactive DLCI—Number of inactive DLCIs. | |
| Input rate | (Redundant LSQ) Rate of bits and packets received on the interface. | None specified |
| Output rate | (Redundant LSQ) Rate of bits and packets transmitted on the interface. | None specified |

Table 63: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|-------------------------|
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Traffic statistics | Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. | detail extensive |
| Frame exceptions | <p>Information about framing exceptions. Includes events recorded under Exception Events for each logical interface.</p> <ul style="list-style-type: none"> • Oversized frames—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb (kilobits). • Errored input frames—Number of input frame errors. • Input on disabled link/bundle—Number of frames received on disabled links. These frames can result either from an inconsistent configuration, or from a bundle or link being brought up or down with traffic actively flowing through it. • Output for disabled link/bundle—Number of frames sent for a disabled or unavailable link. These frames can result either from an inconsistent configuration, or from a bundle being brought up or down while traffic is flowing through it. • Queuing drops—Total number of packets dropped before traffic enters the link services IQ interface. Indicates that the interface is becoming oversubscribed. | extensive |
| Buffering exceptions | <p>Information about buffering exceptions. Includes events recorded under Exception Events for each logical interface:</p> <ul style="list-style-type: none"> • Packet data buffer overflow—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • Fragment data buffer overflow—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. Check the logical interface exception event counters to determine which bundle is responsible. | extensive |

Table 63: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|---|------------------------------|
| Assembly exceptions | <p>(Multilink Frame Relay end-to-end only) Information about assembly exceptions. Includes events recorded under Exception Events for each logical interface.</p> <p>An assembly exception does not necessarily indicate an operational problem with the physical link services IQ interface itself. If multilink-encapsulated traffic is dropped or reordered after a sequence number has been assigned, the interface records one or more exception events. The physical interface can drop multilink-encapsulated fragments itself as a result. Any multilink packets or fragments dropped by the interface itself result in packet or fragment drop counts on individual logical interfaces. If the logical interface drop counts are zero, but exception events are seen, the most likely cause is a problem with the individual link interfaces. Even if the logical interface fragment drop counts are nonzero, excess differential delay or traffic losses on individual interfaces can be the root cause.</p> <ul style="list-style-type: none"> • Fragment timeout—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • Missing sequence number—A gap was detected in the sequence numbers of fragments on a bundle. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • Out-of-order sequence number—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the link services IQ interface unable to correctly process the resulting stream. Check the logical interface exception event counters to determine which bundle is responsible. • Out-of-range sequence number—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. Check the logical interface exception event counters to determine which bundle is responsible. | extensive |
| Hardware errors (sticky) | <p>(Multilink Frame Relay end-to-end only) Information about hardware errors:</p> <ul style="list-style-type: none"> • Data memory error—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. • Control memory error—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. | extensive |
| Egress queues | Total number of egress queues supported on the specified interface. | detail extensive none |

Table 63: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------|---|------------------------------|
| Queue counters | Queue number and its associated user-configured forwarding class name. <ul style="list-style-type: none"> Queued packets—Number of queued packets. Transmitted packets—Number of transmitted packets. Dropped packets—Number of packets dropped by the ASIC's RED mechanism. | detail extensive none |
| Logical Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| Flags | Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Encapsulation | Encapsulation being used: PPP or Multilink PPP. | All levels |
| Bandwidth | Speed at which the interface is running. | All levels |
| Bundle options | (Multilink Frame Relay end-to-end interfaces only) <ul style="list-style-type: none"> MRRU—Configured size of the maximum received reconstructed unit (MRRU): 1500 through 4500 bytes. The default is 1504 bytes. Drop timer period—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: 0 through 2000 milliseconds. Values under 5 ms are not recommended. The default setting is 0, which disables the timer. Sequence number format—Short sequence number header format (MLPPP only). Fragmentation threshold—Configured fragmentation threshold: 64 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation. Links needed to sustain bundle—Minimum number of links to sustain the bundle: 1 through 8. Multilink classes—Number of multilink classes negotiated. Link layer overhead—Percentage of bundle bandwidth to be set aside for link-layer overhead. | detail extensive none |

Table 63: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|---|------------------------------|
| Bundle status
(MLPPP) or
Multilink class status (Multiclass MLPPP) | Information about bundle status: <ul style="list-style-type: none"> • Remote MRRU—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed. • Received sequence number—Sequence number for received packets. • Transmitted sequence number—Sequence number for transmitted packets. • Packet drops—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • Fragment drops—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully, but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • MRRU exceeded—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • Fragment timeout—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • Missing sequence number—A gap was detected in the sequence numbers of fragments on a bundle. • Out-of-order sequence number—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • Out-of-range sequence number—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • Packet data buffer overflow—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • Fragment data buffer overflow—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. | detail extensive none |

Table 63: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|---|------------------------------|
| Statistics | <p>Information about fragments and packets received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).</p> <ul style="list-style-type: none"> • Bundle—Information for each active bundle link. <ul style="list-style-type: none"> • Fragments: Input and Output—Total number and rate of fragments received and transmitted. • Packets: Input and Output—Total number and rate of packets received and transmitted. • Multilink class—(Multiclass MLPPP only) Information about multiclass links used in the multilink operation. • Link—Information about links used in the multilink operation. <ul style="list-style-type: none"> • Link name—Interface name of the link services IQ channel and state information (physical link up or down). • Input and Output—Total number and rate of fragments and packets received and transmitted. | detail extensive |
| NCP state | <p>(PPP) Network Control Protocol state.</p> <ul style="list-style-type: none"> • Conf-ack-received—Acknowledgement was received. • Conf-ack-sent—Acknowledgement was sent. • Conf-req-sent—Request was sent. • Down—NCP negotiation is incomplete (not yet completed or has failed). • Not-configured—NCP is not configured on the interface. • Opened—NCP negotiation is successful. | detail extensive none |
| Protocol | Protocol family configured on the logical interface. | detail extensive none |
| MTU | MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked Adjusted . | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route Table | Routing table in which this address exists. For example, Route table:0 refers to inet.0. | detail extensive |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Addresses, Flags | Information about the addresses configured on the logical interface. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Destination | IP address of the remote side of the connection. | detail extensive none |

Table 63: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|---|------------------------------|
| Local | IP address of the logical interface. | detail extensive none |
| Broadcast | Broadcast address on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| MLPPP Bundle Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| Flags | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| SNMP-Traps | SNMP trap notifications are enabled. | All levels |
| Encapsulation | Encapsulation being used: PPP, Multilink PPP, or Multilink-FR. | All levels |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour:minute:second timezone (hour:minute:second ago)</i> . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) . | detail extensive none |
| Bandwidth | Speed at which the interface is running. | All levels |
| Bundle links information | Information about the bundled links. <ul style="list-style-type: none"> • Active bundle links—Number of active links. • Removed bundle links—Information about links used in the multilink operation. • Disabled bundle links—Number of disabled links. | detail extensive none |

Table 63: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------|--|------------------------------|
| Bundle options | <p>(Multilink Frame Relay end-to-end interfaces only)</p> <ul style="list-style-type: none"> • MRRU—Configured size of the maximum received reconstructed unit (MRRU): 1500 through 4500 bytes. The default is 1504 bytes. • Drop timer period—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: 0 through 2000 milliseconds. Values under 5 ms are not recommended. The default setting is 0, which disables the timer. • Inner PPP Protocol field compression—Inner PPP protocol compression is enabled or disabled. • Sequence number format—Short sequence number header format (MLPPP only). • Fragmentation threshold—Configured fragmentation threshold: 64 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation. • Links needed to sustain bundle—Minimum number of links to sustain the bundle: 1 through 8. • Multilink classes—Number of multilink classes negotiated. • Link layer overhead—Percentage of bundle bandwidth to be set aside for link-layer overhead. | detail extensive none |

Table 63: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|--|------------------------------|
| Bundle status
(MLPPP) | <p>Information about bundle status:</p> <ul style="list-style-type: none"> • Received sequence number—Sequence number for received packets. • Transmit sequence number—Sequence number for transmitted packets. • Packet drops—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • Fragment drops—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • MRRU exceeded—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • Fragment timeout—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • Missing sequence number—A gap was detected in the sequence numbers of fragments on a bundle. • Out-of-order sequence number—Two frames with out-of-order sequence numbers occurred within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • Out-of-range sequence number—A frame was received with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • Packet data buffer overflow—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • Fragment data buffer overflow—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. | detail extensive none |

Table 63: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------|--|------------------|
| Statistics | <p>Information about frames, bytes, and bits per second received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).</p> <p>The bundle, multilink, and network statistics are reported by the Packet Forwarding Engine (PFE). The Multi Link Detail statistics like fragments, non-fragments and LFI are reported by the PIC.</p> <p>However, the PFE reports an extra overhead of 2 bytes in the output when compared with the Multilink Detail Statistics. This is due to the service-cookie in the PFE which does the link demux for the ML header.</p> <p>The difference in the bytes received and transmitted from Network and Multilink interfaces and Multilink statistics for each member link is divided between the ML and the PPP headers. For example the header counter for a long sequence configuration would be as follows.</p> <ul style="list-style-type: none"> • Input side - Total overhead = 6 bytes. <ul style="list-style-type: none"> • ML: 4 bytes of ML header = 1 byte of Flag + 3 bytes of long sequence number. • PPP: 2 bytes of protocol field. • Output side - Total overhead = 11 bytes. <ul style="list-style-type: none"> • ML: 4 bytes of ML Header = 1 byte of Flag + 3 bytes of Long sequence number. • PPP: 5 bytes = 4 bytes of header + 1 byte of Idle flag. • 2 bytes of Service Cookie. • Bundle—Information for each active bundle link. <ul style="list-style-type: none"> • Multilink: Input and Output—Total number and rate of multilink frames, bytes, and bits per second received and transmitted. It is a module connecting LSQ PIC and its member link. Multilink Input displays L2 fragments received from the member link to the LSQ PIC. Multilink Output displays the L2 fragments transmitted from LSQ PIC to the member links. • Network: Input and Output—Total number of network frames, bytes, and bits per second received and transmitted. It refers to the packets transmitted from an ingress interface to the PFE and then to the LSQ PIC. Network Input displays the L3 packets received from the LSQ PIC to the PFE. Network Output displays the L3 packets transmitted from PFE to LSQ PIC. • Link—Information about links used in the multilink operation. <ul style="list-style-type: none"> • Link name—The interface name of the link services IQ channel and state information (physical link <i>up</i> or <i>down</i>) and up time. • Input and Output—Total number and rate of frames, bytes, and bits per second received and transmitted. | extensive |

Table 63: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------------------|---|------------------------------|
| Multilink detail statistics | <p>Frames, bytes, and bits per second received and sent by the bundle. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).</p> <p>The difference in the bytes received and transmitted from the bundle is divided between the ML and the PPP headers. For example the header counter for a long sequence configuration would be as follows:</p> <ul style="list-style-type: none"> • Input side - Total overhead = 6 bytes. <ul style="list-style-type: none"> • ML: 4 bytes of ML header = 1 byte of Flag + 3 bytes of long sequence number. • PPP: 2 bytes of protocol field. • Output side - Total overhead = 9 bytes. <ul style="list-style-type: none"> • ML: 4 bytes of ML Header = 1 byte of Flag + 3 bytes of Long sequence number. • PPP: 5 bytes = 4 bytes of header + 1 byte of Idle flag. • Bundle—Information for the bundle link. <ul style="list-style-type: none"> • Fragments: Input and Output—Total number and rate of multilink fragments received and transmitted. • Non-fragments: Input and Output—Total number and rate of nonfragmented multilink frames received and transmitted. • LFI: Input and Output—Total number and rate of link fragmented and interleaved frames and bytes. | extensive |
| Protocol | Protocol family configured on the logical interface. | detail extensive none |
| MTU | MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked <i>Adjusted</i> . | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route Table | Routing table in which this address exists. For example, Route table:0 refers to inet.0. | detail extensive |
| Addresses, Flags | Information about the addresses configured on the logical interface. Possible values are described in the "Addresses Flags" section under <i>Common Output Fields Description</i> . | detail extensive none |
| Destination | IP address of the remote side of the connection. | detail extensive none |
| Local | IP address of the logical interface. | detail extensive none |
| Broadcast | Broadcast address on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |

Sample Output

show interfaces extensive (MLPPP on Link Services IQ)

```

user@host> show interfaces lsq-0/2/0 extensive
Physical interface: lsq-0/2/0, Enabled, Physical link is Up
  Interface index: 140, SNMP ifIndex: 25, Generation: 23
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2005-06-02 08:54:36 PDT (00:05:45 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :           8872424           229080 bps
    Output bytes  :           9856960           234448 bps
    Input packets :           38202           117 pps
    Output packets:           39453           117 pps
  Frame exceptions:
    Oversized frames           0
    Errored input frames       0
    Input on disabled link/bundle 0
    Output for disabled link/bundle 0
    Queuing drops              0
  Buffering exceptions:
    Packet data buffer overflow 0
    Fragment data buffer overflow 0
  Assembly exceptions:
    Fragment timeout           0
    Missing sequence number     0
    Out-of-order sequence number 0
    Out-of-range sequence number 0
  Hardware errors (sticky):
    Data memory error          0
    Control memory error        0
  Queue counters:

```

| | Queued packets | Transmitted packets | Dropped packets |
|------|----------------|---------------------|-----------------|
| 0 be | 0 | 0 | 0 |
| 1 ef | 0 | 0 | 0 |
| 2 af | 0 | 0 | 0 |
| 3 nc | 0 | 0 | 0 |

```

  Logical interface lsq-0/2/0.0 (Index 66) (SNMP ifIndex 26) (Generation 5)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
  Bandwidth: 256kbps
  Bundle options:
    MRRU           1504
    Drop timer period 2000
    Sequence number format long (24 bits)
    Fragmentation threshold 0
    Links needed to sustain bundle 1
    Multilink classes 0
    Link layer overhead 4.0 %
  Bundle status:
    Remote MRRU           1500
    Received sequence number 0x0
    Transmit sequence number 0x0
    Packet drops           0 (0 bytes)
    Fragment drops         9 (1401 bytes)

```

```

MRRU exceeded          0
Fragment timeout        0
Missing sequence number 0
Out-of-order sequence number 4
Out-of-range sequence number 0
Packet data buffer overflow 0
Fragment data buffer overflow 0
Statistics              Frames    fps          Bytes        bps
Bundle:
Multilink:
  Input :               79827      239          9593009       232288
  Output:              77533      234          9811743       238056
Network:
  Input :               38202      117          8872424       229080
  Output:              39453      117          9856960       234448
Link:
ds-1/0/2:1:1.0 <-- up
  Input :               1114         87          180183        113608
  Output:              1577        118          199215        119064
ds-1/0/2:1:2.0 <-- down
  Input :               1941        152          187948        118680
  Output:              1574        116          199494        118992
Protocol inet, MTU: 1500 [Adjusted]
Flags: User-MTU, MTU-Protocol-Adjusted
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.74.11/24, Local: 10.74.11.10
Protocol iso, MTU: 1500 [Adjusted]
Flags: User-MTU, MTU-Protocol-Adjusted
Protocol mpls, MTU: 1488 [Adjusted], Maximum labels: 3
Flags: User-MTU, MTU-Protocol-Adjusted

```

show interfaces extensive (Multiclass MLPPP on Link Services IQ)

```

user@host> show interfaces extensive lsq-0/2/0
Physical interface: lsq-0/2/0, Enabled, Physical link is Up
Interface index: 140, SNMP ifIndex: 25, Generation: 23
Link-level type: LinkService, MTU: 1504
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps
Last flapped   : 2005-06-02 08:54:36 PDT (00:02:25 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   :          3474024          223704 bps
Output bytes  :          4193992          233888 bps
Input packets :          15809           116 pps
Output packets:          16788           117 pps
Frame exceptions:
Oversized frames          0
Errored input frames      0
Input on disabled link/bundle 0
Output for disabled link/bundle 0
Queuing drops             0
Buffering exceptions:
Packet data buffer overflow 0
Fragment data buffer overflow 0
Assembly exceptions:
Fragment timeout          0
Missing sequence number   0
Out-of-order sequence number 0
Out-of-range sequence number 0
Hardware errors (sticky):

```

| | | | |
|----------------------|----------------|---------------------|-----------------|
| Data memory error | 0 | | |
| Control memory error | 0 | | |
| Queue counters: | Queued packets | Transmitted packets | Dropped packets |
| 0 be | 0 | 0 | 0 |
| 1 ef | 0 | 0 | 0 |
| 2 af | 0 | 0 | 0 |
| 3 nc | 0 | 0 | 0 |

Logical interface lsq-0/2/0.0 (Index 66) (SNMP ifIndex 26) (Generation 5)

Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP

Bandwidth: 256kbps

Bundle options:

| | |
|--------------------------------|----------------|
| MRRU | 1504 |
| Drop timer period | 2000 |
| Sequence number format | long (24 bits) |
| Fragmentation threshold | 0 |
| Links needed to sustain bundle | 1 |
| Multilink classes | 2 |
| Link layer overhead | 4.0 % |

Multilink class 0 status:

| | |
|-------------------------------|---------------------|
| Received sequence number | 0x4c38 |
| Transmit sequence number | 0x4890 |
| Packet drops | 0 (0 bytes) |
| Fragment drops | 2551 (397084 bytes) |
| MRRU exceeded | 0 |
| Fragment timeout | 52 |
| Missing sequence number | 0 |
| Out-of-order sequence number | 953 |
| Out-of-range sequence number | 0 |
| Packet data buffer overflow | 0 |
| Fragment data buffer overflow | 0 |

Multilink class 1 status:

| | |
|-------------------------------|-------------|
| Received sequence number | 0xffffffff |
| Transmit sequence number | 0x3710 |
| Packet drops | 0 (0 bytes) |
| Fragment drops | 0 (0 bytes) |
| MRRU exceeded | 0 |
| Fragment timeout | 0 |
| Missing sequence number | 0 |
| Out-of-order sequence number | 0 |
| Out-of-range sequence number | 0 |
| Packet data buffer overflow | 0 |
| Fragment data buffer overflow | 0 |

| | | | | |
|------------|--------|-----|-------|-----|
| Statistics | Frames | fps | Bytes | bps |
|------------|--------|-----|-------|-----|

Bundle:

Fragments:

| | | | | |
|---------|-------|-----|---------|--------|
| Input : | 33719 | 239 | 4041763 | 231632 |
| Output: | 32371 | 234 | 4096545 | 237488 |

Packets:

| | | | | |
|---------|-------|-----|---------|--------|
| Input : | 15809 | 116 | 3474024 | 223704 |
| Output: | 16788 | 117 | 4193992 | 233888 |

Multilink class 0:

Fragments:

| | | | | |
|---------|-------|---|---|---|
| Input : | 19331 | 0 | 0 | 0 |
| Output: | 0 | 0 | 0 | 0 |

Packets:

| | | | | |
|---------|------|---|---|---|
| Input : | 2064 | 0 | 0 | 0 |
|---------|------|---|---|---|


```

      Output:          1864          0          0          0
Multilink class 1:
  Fragments:
    Input :           0          0          0          0
    Output:         14096          0          0          0
  Packets:
    Input :         14096          0          0          0
    Output:           0          0          0          0
Link:
  ds-1/0/2:1:1.0, Enabled, Physical link is Up
    Input :          20972         151        2030595        118080
    Output:         16184         116        2048468        118488
  ds-1/0/2:1:2.0, Enabled, Physical link is Up
    Input :          12747          88        2011168        113552
    Output:         16187         118        2048077        119000
Protocol inet, MTU: 1500 [Adjusted], Generation: 14, Route table: 0
Flags: User-MTU, MTU-Protocol-Adjusted
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.0.1.0/30, Local: 10.0.1.2, Broadcast: Unspecified,
  Generation: 18

```

show interfaces extensive (MLPPP on Link Services IQ Bundle)

```

user@host> show interfaces lsq-7/1/0.0 extensive
Logical interface lsq-7/1/0.0 (Index 88) (SNMP ifIndex 114) (Generation 188)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-FR
Last flapped: Never
Bandwidth: 256kbps
Bundle links information:
  Active bundle links      2
  Removed bundle links    0
  Disabled bundle links    0
Bundle options:
  MRRU                      1504
  Drop timer period         1500
  Inner PPP Protocol field compression enabled
  Sequence number format    short (12 bits)
  Fragmentation threshold   0
  Links needed to sustain bundle 1
  Multilink classes         0
  Link layer overhead       4.0 %
Bundle status:
  Received sequence number   0xb74
  Transmit sequence number   0xb74
  Packet drops               0 (0 bytes)
  Fragment drops             0 (0 bytes)
  MRRU exceeded              0
  Fragment timeout           0
  Missing sequence number    0
  Out-of-order sequence number 0
  Out-of-range sequence number 0
  Packet data buffer overflow 0
  Fragment data buffer overflow 0
Statistics      Frames      fps      Bytes      bps
Bundle:
Multilink:
  Input :        315381        0      42757818        0
  Output:        315381        0      43388580        0
Network:
  Input :        315381        0      40952064        0
  Output:        315381        0      40952064        0

```

```

Link:
  ds-6/0/0:1:1.0
    Up time: Up since boot
    Input :      63794      0      25146728      0
    Output:      63778      0      25273164      0
  ds-6/0/0:1:2.0
    Up time: Up since boot
    Input :      251587      0      17611090      0
    Output:      251603      0      18115416      0
Multilink detail statistics:
Bundle:
  Fragments:
    Input :      0      0      0      0
    Output:      0      0      0      0
  Non-fragments:
    Input :      293748      0      19387368      0
    Output:      293748      0      20562360      0
  LFI:
    Input :      21633      0      22152192      0
    Output:      21633      0      22325256      0
Protocol inet, MTU: 1500, Generation: 204, Route table: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.0.1.0/30, Local: 10.0.1.2, Broadcast:
Unspecified, Generation: 214

```

show interfaces extensive (MFR on Link Services IQ Bundle)

```

user@host> show interfaces lsq-1/0/0:0 extensive
Physical interface: lsq-1/0/0:0, Enabled, Physical link is Up
Interface index: 179, SNMP ifIndex: 746, Generation: 182
Link-level type: Multilink-FR-UNI-NNI, MTU: 1508
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
Last flapped   : 2010-11-15 01:11:00 PST (00:31:58 ago)
Statistics last cleared: Never
Hold-times     : Up 0 ms, Down 0 ms
Multilink Frame Relay UNI NNI bundle options:
  Device type      DCE
  MRRU             1508
  Bandwidth        1536kbps
  Fragmentation threshold 0
  Red differential delay limit 120
  Yellow differential delay limit 72
  Red differential delay action Remove link
  Reassembly drop timer 65535
  Links needed to sustain bundle 1
  Link layer overhead 4.0 %
  LIP Hello timer 10
    Acknowledgement timer 4
    Acknowledgement retries 2
  Bundle class     A
  LMI type         Consortium
    T391 LIV polling timer 10
    T392 polling verification timer 15
    N391 full status polling count 6
    N392 error threshold 3
    N393 monitored event count 4
  Consortium LMI settings: n392dce 3, n393dce 4, t392dce 15 seconds
LMI statistics:
  Input : 188 (last seen 00:00:01 ago)

```

```

Output: 189 (last sent 00:00:01 ago)
DTE statistics:
  Enquiries sent : 0
  Full enquiries sent : 0
  Enquiry responses received : 0
  Full enquiry responses received : 0
DCE statistics:
  Enquiries received : 157
  Full enquiries received : 31
  Enquiry responses sent : 158
  Full enquiry responses sent : 31
Common statistics:
  Unknown messages received : 0
  Asynchronous updates received : 0
  Out-of-sequence packets received : 0
  Keepalive responses timedout : 0
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Multilink Frame Relay UNI NNI bundle errors:
  Packet drops 0 (0 bytes)
  Fragment drops 0 (0 bytes)
  MRRU exceeded 0
  Exception events 0
Multilink Frame Relay UNI NNI bundle statistics:
      Frames      fps      Bytes      bps

Multilink:
  Input : 0 0 0 0
  Output: 0 0 0 0
Network:
  Input : 0 0 0 0
  Output: 0 0 0 0
Multilink Frame Relay UNI NNI bundle links information:
  Active bundle links 1
  Removed bundle links 0
  Disabled bundle links 0
Multilink Frame Relay UNI NNI active bundle links statistics:
      Frames      fps      Bytes      bps

t1-7/0/0:1:3.0
Up time: 00:31:24
  Input : 0 0 0 0
  Output: 0 0 0 0
  Current differential delay 0.0 ms
  Recent high differential delay 0.0 ms
  Times over red diff delay 0
  Times over yellow diff delay 0
LIP:add_lnk lnk_ack lnk_rej hello hel_ack lnk_rem rem_ack
Rcv: 2 2 0 0 189 0 0
Xmt: 2 1 0 189 0 0 0

```

Logical interface lsq-1/0/0:2.0 (Index 77) (SNMP ifIndex 751) (Generation 142)

Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI

```

Last flapped: 2010-11-15 01:11:40 PST (00:31:18 ago)
Bundle status:
  Received sequence number      0xfff
  Transmit sequence number      0x0
  Packet drops                  0 (0 bytes)
  Fragment drops                0 (0 bytes)
  MRRU exceeded                 0
  Fragment timeout              0
  Missing sequence number       0
  Out-of-order sequence number  0
  Out-of-range sequence number  0
  Packet data buffer overflow    0
  Fragment data buffer overflow  0
Statistics      Frames      fps      Bytes      bps
Bundle:
Multilink:
  Input :        0          0          0          0
  Output:        0          0          0          0
Network:
  Input :        0          0          0          0
  Output:        0          0          0          0
Link:
  t1-7/0/0:1:3.0
  Up time: 00:31:24
  Input :        0          0          0          0
  Output:        0          0          0          0
Multilink detail statistics:
Bundle:
Fragments:
  Input :        0          0          0          0
  Output:        0          0          0          0
Non-fragments:
  Input :        0          0          0          0
  Output:        0          0          0          0
Protocol inet, MTU: 1500, Generation: 153, Route table: 0
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.0.1.8/30, Local: 10.0.1.9, Broadcast: Unspecified,
Generation: 154
DLCI 12
Flags: Active
Total down time: 00:00:32 sec, Last down: 00:31:50 ago
Traffic statistics:
  Input bytes :          0
  Output bytes :          0
  Input packets:         0
  Output packets:        0
DLCI statistics:
  Active DLCI :1 Inactive DLCI :0

```

show interfaces (Multiclass MLPPP on Link Services IQ)

```

user@host> show interfaces extensive lsq-0/2/0
Physical interface: lsq-0/2/0, Enabled, Physical link is Up
Interface index: 140, SNMP ifIndex: 25, Generation: 23
Link-level type: LinkService, MTU: 1504
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps
Last flapped : 2005-06-02 08:54:36 PDT (00:02:25 ago)
Statistics last cleared: Never
Traffic statistics:

```

```

Input bytes :          3474024          223704 bps
Output bytes :         4193992          233888 bps
Input packets:         15809           116 pps
Output packets:        16788           117 pps
Frame exceptions:
  Oversized frames      0
  Errored input frames  0
  Input on disabled link/bundle 0
  Output for disabled link/bundle 0
  Queuing drops        0
Buffering exceptions:
  Packet data buffer overflow 0
  Fragment data buffer overflow 0
Assembly exceptions:
  Fragment timeout      0
  Missing sequence number 0
  Out-of-order sequence number 0
  Out-of-range sequence number 0
Hardware errors (sticky):
  Data memory error     0
  Control memory error  0
Queue counters:         Queued packets  Transmitted packets  Dropped packets

0 be                    0                0                0
1 ef                    0                0                0
2 af                    0                0                0
3 nc                    0                0                0

```

Logical interface lsq-0/2/0.0 (Index 66) (SNMP ifIndex 26) (Generation 5)

Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP

Bandwidth: 256kbps

Bundle options:

```

MRRU                    1504
Drop timer period      2000
Sequence number format long (24 bits)
Fragmentation threshold 0
Links needed to sustain bundle 1
Multilink classes      2
Link layer overhead    4.0 %

```

Multilink class 0 status:

```

Received sequence number 0x4c38
Transmit sequence number 0x4890
Packet drops             0 (0 bytes)
Fragment drops           2551 (397084 bytes)
MRRU exceeded            0
Fragment timeout         52
Missing sequence number  0
Out-of-order sequence number 953
Out-of-range sequence number 0
Packet data buffer overflow 0
Fragment data buffer overflow 0

```

Multilink class 1 status:

```

Received sequence number 0xffffffff
Transmit sequence number 0x3710
Packet drops             0 (0 bytes)
Fragment drops           0 (0 bytes)
MRRU exceeded            0
Fragment timeout         0

```

```

Missing sequence number      0
Out-of-order sequence number 0
Out-of-range sequence number 0
Packet data buffer overflow  0
Fragment data buffer overflow 0
Statistics      Frames      fps      Bytes      bps
Bundle:
Fragments:
  Input :      33719      239      4041763      231632
  Output:      32371      234      4096545      237488
Packets:
  Input :      15809      116      3474024      223704
  Output:      16788      117      4193992      233888
Multilink class 0:
Fragments:
  Input :      19331      0      0      0
  Output:      0      0      0      0
Packets:
  Input :      2064      0      0      0
  Output:      1864      0      0      0
Multilink class 1:
Fragments:
  Input :      0      0      0      0
  Output:      14096      0      0      0
Packets:
  Input :      14096      0      0      0
  Output:      0      0      0      0
Link:
ds-1/0/2:1:1.0, Enabled, Physical link is Up
  Input :      20972      151      2030595      118080
  Output:      16184      116      2048468      118488
ds-1/0/2:1:2.0, Enabled, Physical link is Up
  Input :      12747      88      2011168      113552
  Output:      16187      118      2048077      119000
Protocol inet, MTU: 1500 [Adjusted], Generation: 14, Route table: 0
Flags: User-MTU, MTU-Protocol-Adjusted
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.0.1.0/30, Local: 10.0.1.2, Broadcast: Unspecified,
  Generation: 18

```

show interfaces (Redundant Adaptive Services)

| | |
|---------------------------------|---|
| Syntax | <pre>show interfaces <i>rspnumber</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M Series and T Series routers only) Display status information about the specified redundant adaptive services configuration. |
| Options | <p><i>rspnumber</i>—Display standard status information about the specified redundant adaptive services configuration.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| List of Sample Output | show interfaces extensive (Redundant Adaptive Services) on page 1909 |
| Output Fields | See the output field table for the show interfaces (Adaptive Services) command. |

Sample Output

show interfaces extensive (Redundant Adaptive Services)

```
user@host> show interfaces rsp0 extensive
Physical interface: rsp0, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 40, Generation: 44
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Redundancy-Device 16384
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times    : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped  : 2005-03-11 18:36:37 UTC (00:00:08 ago)
  Statistics last cleared: Never
  Traffic statistics:
```

```
Input bytes :                0                0 bps
Output bytes :                0                0 bps
Input packets:               0                0 pps
Output packets:              0                0 pps
```

Input errors:

```
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
Policed discards: 0, Resource errors: 0
```

Output errors:

```
Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
Resource errors: 0
```

Logical interface rsp0.0 (Index 68) (SNMP ifIndex 42) (Generation 30)

Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services

Traffic statistics:

```
Input bytes :                0
Output bytes :                0
Input packets:               0
Output packets:              0
```

Local statistics:

```
Input bytes :                0
Output bytes :                0
Input packets:               0
Output packets:              0
```

Transit statistics:

```
Input bytes :                0                0 bps
Output bytes :                0                0 bps
Input packets:               0                0 pps
Output packets:              0                0 pps
```

Protocol inet, MTU: 9192, Generation: 37, Route table: 0

Flags: Receive-options, Receive-TTL-Exceeded

show interfaces (Redundant Link Services IQ)

| | |
|---------------------------------|--|
| Syntax | <pre>show interfaces rlsqnumber <brief detail extensive terse> <descriptions> <media> <queue> <routing> <snmp-index snmp-index> <statistics></pre> |
| Release Information | Command introduced in Junos OS Release 7.6. |
| Description | (M Series and T Series routers only) Display status information about the specified redundant link services intelligent queuing (IQ) configuration. |
| Options | <p>rlsqnumber—Redundant link services IQ interface name. The logical interface number range of values is 0 through 127.</p> <p>none—Display standard status information about the specified redundant link services IQ configuration.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>queue—(Optional) Display queue information about network interfaces.</p> <p>routing—(Optional) Display routing information about network interfaces.</p> <p>snmp-index snmp-index—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show interfaces (Redundant Link Services IQ) on page 1922</p> <p>show interfaces brief (Redundant Link Services IQ) on page 1922</p> <p>show interfaces detail (Redundant Link Services IQ) on page 1923</p> <p>show interfaces extensive (Redundant Link Services IQ) on page 1924</p> |
| Output Fields | Table 64 on page 1911 lists the output fields for the show interfaces (redundant link services IQ) command. Output fields are listed in the approximate order in which they appear. |

Table 64: show interfaces (Redundant Link Services IQ) Output Fields

| Field Name | Field Description | Level of Output |
|--------------------|-------------------|-----------------|
| Physical Interface | | |

Table 64: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|--|------------------------------|
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface's index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| Link-level type | Encapsulation being used on the physical interface:
Multilink-Frame-Relay-UNI-NNI (default), LinkService , Frame-relay , Frame-relay-ccc , or Frame-relay-tcc . | All levels |
| MTU | Maximum transmission unit size on the physical interface. | All levels |
| Device flags | Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) . | detail extensive none |
| Input rate | (Redundant LSQ) Rate of bits and packets received on the interface. | None specified |
| Output rate | (Redundant LSQ) Rate of bits and packets transmitted on the interface. | None specified |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Traffic statistics | Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. | detail extensive |

Table 64: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------------|---|------------------|
| Frame exceptions | <p>Information about framing exceptions. Includes events recorded under Exception Events for each logical interface.</p> <ul style="list-style-type: none"> • Oversized frames—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb (kilobits). • Errored input frames—Number of input frame errors. • Input on disabled link/bundle—Number of frames received on disabled links. These frames can result either from an inconsistent configuration, or from a bundle or link being brought up or down with traffic actively flowing through it. • Output for disabled link/bundle—Number of frames sent for a disabled or unavailable link. These frames can result either from an inconsistent configuration, or from a bundle being brought up or down while traffic is flowing through it. • Queuing drops—Total number of packets dropped before traffic enters the link services IQ interface. Indicates that the interface is becoming oversubscribed. | extensive |
| Buffering exceptions | <p>Information about buffering exceptions. Includes events recorded under Exception Events for each logical interface:</p> <ul style="list-style-type: none"> • Packet data buffer overflow—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • Fragment data buffer overflow—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. Check the logical interface exception event counters to determine which bundle is responsible. | extensive |

Table 64: show interfaces (Redundant Link Services IQ) Output Fields (continued)

| Field Name | Field Description | Level of Output |
|---------------------------------|---|------------------------------|
| Assembly exceptions | <p>(Multilink Frame Relay end-to-end only) Information about assembly exceptions. Includes events recorded under Exception Events for each logical interface.</p> <p>An assembly exception does not necessarily indicate an operational problem with the physical link services IQ interface itself. If multilink-encapsulated traffic is dropped or reordered after a sequence number has been assigned, the interface records one or more exception events. The physical interface can drop multilink-encapsulated fragments itself as a result. Any multilink packets or fragments dropped by the interface itself result in packet or fragment drop counts on individual logical interfaces. If the logical interface drop counts are zero, but exception events are seen, the most likely cause is a problem with the individual link interfaces. Even if the logical interface fragment drop counts are nonzero, excess differential delay or traffic losses on individual interfaces can be the root cause.</p> <ul style="list-style-type: none"> • Fragment timeout—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • Missing sequence number—A gap was detected in the sequence numbers of fragments on a bundle. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • Out-of-order sequence number—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the link services IQ interface unable to correctly process the resulting stream. Check the logical interface exception event counters to determine which bundle is responsible. • Out-of-range sequence number—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. Check the logical interface exception event counters to determine which bundle is responsible. | extensive |
| Hardware errors (sticky) | <p>(Multilink Frame Relay end-to-end only) Information about hardware errors:</p> <ul style="list-style-type: none"> • Data memory error—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. • Control memory error—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. | extensive |
| Egress queues | Total number of egress queues supported on the specified interface. | detail extensive none |

Table 64: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------|---|------------------------------|
| Queue counters | Queue number and its associated user-configured forwarding class name. <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. | detail extensive none |
| Logical Interface | | |
| Logical interface | Name of the logical interface | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support | detail extensive |
| Flags | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Encapsulation | Encapsulation being used: PPP or Multilink PPP. | All levels |
| Bandwidth | Speed at which the interface is running. | All levels |
| Bundle options | (Multilink Frame Relay end-to-end interfaces only) <ul style="list-style-type: none"> • MRRU—Configured size of the maximum received reconstructed unit (MRRU): 1500 through 4500 bytes. The default is 1504 bytes. • Drop timer period—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: 0 through 2000 milliseconds. Values under 5 ms are not recommended. The default setting is 0, which disables the timer. • Sequence number format—Short sequence number header format (MLPPP only). • Fragmentation threshold—Configured fragmentation threshold: 64 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation. • Links needed to sustain bundle—Minimum number of links to sustain the bundle: 1 through 8. • Multilink classes—Number of multilink classes negotiated. • Link layer overhead—Percentage of bundle bandwidth to be set aside for link-layer overhead. | detail extensive none |

Table 64: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|--|-----------------------|
| Bundle status
(MLPPP) or
Multilink class status
(MC-MLPPP) | Information about bundle status: <ul style="list-style-type: none"> • Remote MRRU—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed. • Received sequence number—Sequence number for received packets. • Transmitted sequence number—Sequence number for transmitted packets. • Packet drops—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • Fragment drops—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • MRRU exceeded—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • Fragment timeout—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • Missing sequence number—A gap was detected in the sequence numbers of fragments on a bundle. • Out-of-order sequence number—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • Out-of-range sequence number—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • Packet data buffer overflow—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • Fragment data buffer overflow—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. | detail extensive none |

Table 64: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|---|------------------------------|
| Statistics | <p>Information about fragments and packets received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).</p> <ul style="list-style-type: none"> • Bundle—Information for each active bundle link. <ul style="list-style-type: none"> • Fragments: Input and Output—Total number and rate of fragments received and transmitted. • Packets: Input and Output—Total number and rate of packets received and transmitted. • Multilink class—(MC-MLPPP only) Information about multiclass links used in the multilink operation. • Link—Information about links used in the multilink operation. <ul style="list-style-type: none"> • Link name—Interface name of the link services IQ channel and state information (physical link up or down). • Input and Output—Total number and rate of fragments and packets received and transmitted. | detail extensive |
| NCP state | <p>(PPP) Network Control Protocol state.</p> <ul style="list-style-type: none"> • Conf-ack-received—Acknowledgement was received. • Conf-ack-sent—Acknowledgement was sent. • Conf-req-sent—Request was sent. • Down—NCP negotiation is incomplete (not yet completed or has failed). • Not-configured—NCP is not configured on the interface. • Opened—NCP negotiation is successful. | detail extensive none |
| Protocol | Protocol family configured on the logical interface. | detail extensive none |
| MTU | MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked Adjusted . | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route Table | Routing table in which this address exists. For example, Route table:0 refers to inet.0. | detail extensive |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Addresses, Flags | Information about the addresses configured on the logical interface. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Destination | IP address of the remote side of the connection. | detail extensive none |

Table 64: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|---|------------------------------|
| Local | IP address of the logical interface. | detail extensive none |
| Broadcast | Broadcast address on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| MLPPP Bundle Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| Flags | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| SNMP-Traps | SNMP trap notifications are enabled. | All levels |
| Encapsulation | Encapsulation being used: PPP, Multilink PPP or Multilink-FR. | All levels |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour:minute:second timezone (hour:minute:second ago)</i> . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) . | detail extensive none |
| Bandwidth | Speed at which the interface is running. | All levels |
| Bundle links information | Information about the bundled links. <ul style="list-style-type: none"> • Active bundle links—Number of active links. • Removed bundle links—Information about links used in the multilink operation. • Disabled bundle links—Number of disabled links. | detail extensive none |

Table 64: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------|---|------------------------------|
| Bundle options | <p>(Multilink Frame Relay end-to-end interfaces only)</p> <ul style="list-style-type: none"> • MRRU—Configured size of the maximum received reconstructed unit (MRRU): 1500 through 4500 bytes. The default is 1504 bytes. • Drop timer period—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: 0 through 2000 milliseconds. Values under 5 ms are not recommended. The default setting is 0, which disables the timer. • Inner PPP Protocol field compression—Inner PPP protocol compression is enabled or disabled. • Sequence number format—Short sequence number header format (MLPPP only). • Fragmentation threshold—Configured fragmentation threshold: 64 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation. • Links needed to sustain bundle—Minimum number of links to sustain the bundle: 1 through 8. • Multilink classes—Number of multilink classes negotiated. • Link layer overhead—Percentage of bundle bandwidth to be set aside for link-layer overhead. | detail extensive none |

Table 64: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|--|-----------------------|
| Bundle status
(MLPPP) | <p>Information about bundle status:</p> <ul style="list-style-type: none"> • Received sequence number—Sequence number for received packets. • Transmit sequence number—Sequence number for transmitted packets. • Packet drops—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • Fragment drops—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • MRRU exceeded—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • Fragment timeout—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • Missing sequence number—A gap was detected in the sequence numbers of fragments on a bundle. • Out-of-order sequence number—Two frames with out-of-order sequence numbers occurred within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • Out-of-range sequence number—A frame was received with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • Packet data buffer overflow—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • Fragment data buffer overflow—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. | detail extensive none |

Table 64: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------------------|---|------------------------------|
| Statistics | <p>Information about frames, bytes, and bits per second received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).</p> <ul style="list-style-type: none"> • Bundle—Information for each active bundle link. <ul style="list-style-type: none"> • Multilink: Input and Output—Total number and rate of multilink frames, bytes, and bits per second received and transmitted. • Network: Input and Output—Total number of multilink frames, bytes, and bits per second received and transmitted. • Link—Information about links used in the multilink operation. <ul style="list-style-type: none"> • Link name is the interface name of the link services IQ channel and state information (physical link up or down) and up time. • Input and Output—Total number and rate of frames, bytes, and bits per second received and transmitted. | extensive |
| Multilink detail statistics | <p>Frames, bytes, and bits per second received and sent by the bundle. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).</p> <ul style="list-style-type: none"> • Bundle—Information for the bundle link. <ul style="list-style-type: none"> • Fragments: Input and Output—Total number and rate of multilink fragments received and transmitted. • Non-fragments: Input and Output—Total number and rate of nonfragmented multilink frames received and transmitted. • LFI: Input and Output—Total number and rate of link fragmented and interleaved frames and bytes. | extensive |
| Protocol | Protocol family configured on the logical interface. | detail extensive none |
| MTU | MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked Adjusted . | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route Table | Routing table in which this address exists. For example, Route table:0 refers to inet.0. | detail extensive |
| Addresses, Flags | Information about the addresses configured on the logical interface. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Destination | IP address of the remote side of the connection. | detail extensive none |
| Local | IP address of the logical interface. | detail extensive none |
| Broadcast | Broadcast address on the logical interface. | detail extensive none |

Table 64: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------|--|------------------|
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |

Sample Output

show interfaces (Redundant Link Services IQ)

```

user@host> show interfaces rlsq0
Physical interface: rlsq0, Enabled, Physical link is Up
  Interface index: 196, SNMP ifIndex: 27
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Last flapped   : Never
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)

Logical interface rlsq0.0 (Index 72) (SNMP ifIndex 88)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
  Bandwidth: 0
  Statistics          Frames          fps          Bytes          bps
  Bundle:
    Fragments:
      Input :           3             0           255            0
      Output:           3             0           264            0
    Packets:
      Input :           3             0           252            0
      Output:           0             0            0            0
  Link:
    t1-1/3/0:1.0
      Input :           3             0           255            0
      Output:           0             0            0            0
    t1-1/3/0:2.0
      Input :           0             0            0            0
      Output:           3             0           264            0
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  Protocol inet, MTU: 1500
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
  Destination: 2.2.2.0/30, Local: 2.2.2.1

```

show interfaces brief (Redundant Link Services IQ)

```

user@host> show interfaces rlsq0 brief
Physical interface: rlsq0, Enabled, Physical link is Up
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000

Logical interface rlsq0.0
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
  inet 2.2.2.1/30

```

show interfaces detail (Redundant Link Services IQ)

```

user@host> show interfaces rlsq0 detail
Physical interface: rlsq0, Enabled, Physical link is Up
  Interface index: 196, SNMP ifIndex: 27, Generation: 144
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :          252          0 bps
    Output bytes :          276          0 bps
    Input packets:           3          0 pps
    Output packets:          3          0 pps
  Frame exceptions:
    Oversized frames          0
    Errored input frames      0
    Input on disabled link/bundle 0
    Output for disabled link/bundle 0
    Queuing drops            0
  Buffering exceptions:
    Packet data buffer overflow 0
    Fragment data buffer overflow 0
  Assembly exceptions:
    Fragment timeout          0
    Missing sequence number    0
    Out-of-order sequence number 0
    Out-of-range sequence number 0
  Hardware errors (sticky):
    Data memory error          0
    Control memory error        0
  Egress queues: 8 supported, 4 in use
  Queue counters:

```

| | Queued packets | Transmitted packets | Dropped packets |
|----------------|----------------|---------------------|-----------------|
| 0 be | 0 | 0 | 0 |
| 1 expedited-fo | 0 | 0 | 0 |
| 2 assured-forw | 0 | 0 | 0 |
| 3 network-cont | 0 | 0 | 0 |

```

Logical interface rlsq0.0 (Index 72) (SNMP ifIndex 88) (Generation 31)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
  Bandwidth: 0
  Bundle options:
    MRRU          1504
    Remote MRRU    N/A
    Drop timer period 2000
    Sequence number format long (24 bits)
    Fragmentation threshold 0
    Links needed to sustain bundle 1
    Multilink classes 0
    Link layer overhead 4.0 %
  Bundle status:
    Received sequence number 0xffffffff
    Transmit sequence number 0x0
    Packet drops 0 (0 bytes)
    Fragment drops 0 (0 bytes)

```

```

MRRU exceeded          0
Fragment timeout       0
Missing sequence number 0
Out-of-order sequence number 0
Out-of-range sequence number 0
Packet data buffer overflow 0
Fragment data buffer overflow 0
Statistics              Frames      fps          Bytes      bps
Bundle:
Fragments:
  Input :              3          0          255        0
  Output:              3          0          264        0
Packets:
  Input :              3          0          252        0
  Output:              0          0           0        0
Link:
t1-1/3/0:1.0
  Input :              3          0          255        0
  Output:              0          0           0        0
t1-1/3/0:2.0
  Input :              0          0           0        0
  Output:              3          0          264        0
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
Protocol inet, MTU: 1500, Generation: 43, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 2.2.2.0/30, Local: 2.2.2.1, Broadcast: Unspecified,
Generation: 45

```

[show interfaces extensive \(Redundant Link Services IQ\)](#)

The output for the **show interfaces rlsq extensive** command is identical to that for the **show interfaces rlsq detail** command. For sample output, see [show interfaces detail \(Redundant Link Services IQ\) on page 1923](#).

show interfaces load-balancing

| | |
|---------------------------------|---|
| Syntax | show interfaces load-balancing
<detail> |
| Release Information | Command introduced in Junos OS Release 11.4. |
| Description | Display status information about load balancing on aggregated Multiservices (AMS) interfaces. |
| Options | none —Display standard information about status of all AMS interfaces.
detail —(Optional) Display detailed status of all AMS interfaces. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • Understanding Aggregated Multiservices Interfaces on page 599 • Example: Configuring an Aggregated Multiservices Interface (AMS) on page 608 |
| List of Sample Output | show interfaces load-balancing on page 1927
show interfaces load-balancing detail on page 1927 |
| Output Fields | Table 65 on page 1925 lists the output fields for the show interfaces load-balancing (aggregated Multiservices interfaces) command. Output fields are listed in the approximate order in which they appear. |

Table 65: Aggregated Multiservices show interfaces load-balancing Output Fields

| Field Name | Field Description | Level of Output |
|--------------------------|--|-----------------|
| Interface | Name of the aggregated Multiservices (AMS) interface. | All levels |
| State | Status of AMS interfaces: <ul style="list-style-type: none"> • Up—Interface is configured and operational. • Coming Up—Interface is becoming operational. • Wait Timer—Interface is waiting for member interfaces (mams) to come online. • Members Seen—Member interfaces (mams) are available. • Wait for Members—Member interfaces (mams) are not available. | All levels |
| Last change | Time elapsed since the last change to the interface. Changes that affect the elapsed time displayed include internal events that may not have changed the state of the member | All levels |
| Member count | Number of member PICs (mams) that are part of the aggregated interface. | All levels |
| Members interface | List of all member PICs (mams) that are part of the aggregated interface. | detail |

Table 65: Aggregated Multiservices show interfaces load-balancing Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------|---|-----------------|
| Weight | Weight associated with each member PIC for load balancing. The minimum weight is 1, maximum weight is 100; default weight is 10. | detail |
| State | Status of each member PIC (mams) : <ul style="list-style-type: none">• Invalid—Configured interface is not valid.• Down—Interface is not operational.• Active—Interface is configured and operational.• Discard—Interface has been discarded.• Inactive—Configured interface is not online.• Backup—Interface has been configured as backup. | detail |

Sample Output

show interfaces load-balancing

```
user@host> show interfaces load-balancing
Interface  State      Last change  Member count
ams0       Up         1d 00:50    2
ams1       Up         00:00:59    2
```

show interfaces load-balancing detail

```
user@host> show interfaces load-balancing detail
Load-balancing interfaces detail
Interface      : ams0
State          : Up
Last change    : 1d 00:51
Member count   : 2
Members       :
  Interface    Weight  State
  mams-2/0/0   10     Active
  mams-2/1/0   10     Active
```

show interfaces redundancy


| | |
|---------------------------------|--|
| Syntax | show interfaces redundancy
<brief detail> |
| Release Information | Command introduced before Junos OS Release 7.4.
detail option added in Junos OS Release 10.0. |
| Description | (M Series, T Series, and MX Series routers only) Display general information about adaptive services and link services intelligent queuing (IQ) interfaces and aggregated Ethernet interfaces redundancy. |
| | <div>  <p>NOTE: When you run the show interfaces redundancy command on an MX80 router, it displays the error message, error:the redundancy-interface-process subsystem is not running. This is because an MX80 router does not have a redundant FPC and does not support link protection.</p> </div> |
| Options | brief detail —(Optional) Display the specified level of output. |
| Required Privilege Level | view |
| List of Sample Output | show interfaces redundancy on page 1929
show interfaces redundancy (Aggregated Ethernet) on page 1929
show interfaces redundancy detail on page 1929 |
| Output Fields | Table 66 on page 1928 lists the output fields for the show interfaces redundancy command. Output fields are listed in the approximate order in which they appear. |

Table 66: show interfaces redundancy Output Fields

| Field Name | Field Description | Level of Output |
|--------------------|---|-----------------|
| Interface | Name of the redundant adaptive services, link services IQ interfaces, or aggregated Ethernet interfaces. | All levels |
| State | State of the redundant interface: Not present , On primary , On secondary or Waiting for primary MS PIC . | All levels |
| Last Change | <p>Timestamp for the last change in status. This value resets after a master Routing Engine switchover event if any of the following conditions is met:</p> <ul style="list-style-type: none"> • GRES is not configured on the router. • The rlsq interface is configured without the hot-standby or warm-standby statements and the backup lsq interface was active before the switchover. • No logical interfaces are configured or all of the configured logical interfaces are down at the time of the switchover. | All levels |

Table 66: show interfaces redundancy Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------|---|-----------------|
| Primary | Name of the interface configured to be the primary interface. | All levels |
| Secondary | Name of the interface configured to be the backup interface. | All levels |
| Current Status | Physical status of the primary and secondary interfaces. | All levels |
| Mode | Standby mode. | detail |

Sample Output

show interfaces redundancy

```

user@host> show interfaces redundancy
Interface  State           Last change  Primary    Secondary   Current status
rsp0       Not present                    sp-1/0/0   sp-0/2/0   both down
rsp1       On secondary    1d 23:56    sp-1/2/0   sp-0/3/0   primary down
rsp2       On primary      10:10:27    sp-1/3/0   sp-0/2/0   secondary down
rlsq0      On primary      00:06:24    lsq-0/3/0   lsq-1/0/0   both up

```

show interfaces redundancy (Aggregated Ethernet)

```

user@host> show interfaces redundancy
Interface  State           Last change  Primary    Secondary   Current status
rlsq0      On secondary    00:56:12    lsq-4/0/0   lsq-3/0/0   both up

ae0
ae1
ae2
ae3
ae4

```

show interfaces redundancy detail

```

user@host> show interfaces redundancy detail
Interface      : rlsq0
State          : On primary
Last change    : 00:45:47
Primary        : lsq-0/2/0
Secondary      : lsq-1/2/0
Current status : both up
Mode           : hot-standby

Interface      : rlsq0:0
State          : On primary
Last change    : 00:45:46
Primary        : lsq-0/2/0:0
Secondary      : lsq-1/2/0:0
Current status : both up
Mode           : warm-standby

```

show security pki ca-certificate

| | |
|---------------------------------|--|
| Syntax | show security pki ca-certificate
<brief detail>
<ca-profile <i>ca-profile-name</i> > |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | Display information about certificate authority (CA) digital certificates installed in the router. |
| Options | <p>none—(Same as brief) Display information about all CA digital certificates.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>ca-profile <i>ca-profile-name</i>—(Optional) Display information about only the specified CA profile.</p> |
| Required Privilege Level | view |
| List of Sample Output | show security pki ca-certificate on page 1931
show security pki ca-certificate detail on page 1932 |
| Output Fields | Table 67 on page 1930 lists the output fields for the show security pki ca-certificate command. Output fields are listed in the approximate order in which they appear. |

Table 67: show security pki ca-certificate Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------------|---|-------------------|
| Certificate identifier | Name of the digital certificate. | All levels |
| Certificate version | Revision number of the digital certificate. | detail |
| Serial number | Unique serial number of the digital certificate. | detail |
| Issued by | Authority that issued the digital certificate. | none brief |
| Issued to | Device that was issued the digital certificate. | none brief |
| Issuer | <p>Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are:</p> <ul style="list-style-type: none"> • Common name—Name of the authority. • Organization—Organization of origin. • Organizational unit—Department within an organization. • State—State of origin. • Country—Country of origin. | detail |

Table 67: show security pki ca-certificate Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------------|--|-----------------|
| Subject | Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: <ul style="list-style-type: none"> • Common name—Name of the requestor. • Organization—Organization of origin. • Organizational unit—Department within an organization. • State—State of origin. • Country—Country of origin. | detail |
| Validity | Time period when the digital certificate is valid. Values are: <ul style="list-style-type: none"> • Not before—Start time when the digital certificate becomes valid. • Not after—End time when the digital certificate becomes invalid. | All levels |
| Public key algorithm | Encryption algorithm used with the private key, such as rsaEncryption(1024 bits) . | All levels |
| Signature algorithm | Encryption algorithm that the CA used to sign the digital certificate, such as sha1WithRSAEncryption . | detail |
| Fingerprint | Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate. | detail |
| Distribution CRL | Distinguished name information and the URL for the certificate revocation list (CRL) server. | detail |
| Use for key | Use of the public key, such as Certificate signing , CRL signing , Digital signature , or Key encipherment . | detail |

Sample Output

show security pki ca-certificate

```

user@host> show security pki ca-certificate
Certificate identifier: entrust
  Issued to: juniper, Issued by: juniper
  Validity:
    Not before: 2005 Oct 18th, 23:54:22 GMT
    Not after: 2025 Oct 19th, 00:24:22 GMT
  Public key algorithm: rsaEncryption(1024 bits)

Certificate identifier: entrust
  Issued to: First Officer, Issued by: juniper
  Validity:
    Not before: 2005 Oct 18th, 23:55:59 GMT
    Not after: 2008 Oct 19th, 00:25:59 GMT
  Public key algorithm: rsaEncryption(1024 bits)

Certificate identifier: entrust
  Issued to: First Officer, Issued by: juniper
  Validity:
    Not before: 2005 Oct 18th, 23:55:59 GMT

```

Not after: 2008 Oct 19th, 00:25:59 GMT
Public key algorithm: rsaEncryption(1024 bits)

show security pki ca-certificate detail

```

user@host> show security pki ca-certificate detail
Certificate identifier: entrust
Certificate version: 3
Serial number: 4355 9235
Issuer:
  Organization: juniper, Country: us
Subject:
  Organization: juniper, Country: us
Validity:
  Not before: 2005 Oct 18th, 23:54:22 GMT
  Not after: 2025 Oct 19th, 00:24:22 GMT
Public key algorithm: rsaEncryption(1024 bits)
cb:9e:2d:c0:70:f8:ea:3c:f2:b5:f0:02:48:87:dc:68:99:a3:57:4f
0e:b9:98:0b:95:47:0d:1f:97:7c:53:17:dd:1a:f8:da:e5:08:d1:1c
78:68:1f:2f:72:9f:a2:cf:81:e3:ce:c5:56:89:ce:f0:97:93:fa:36
19:3e:18:7d:8c:9d:21:fe:1f:c3:87:8d:b3:5d:f3:03:66:9d:16:a7
bf:18:3f:f0:7a:80:f0:62:50:43:83:4f:0e:d7:c6:42:48:c0:8a:b2
c7:46:30:38:df:9b:dc:bc:b5:08:7a:f3:cd:64:db:2b:71:67:fe:d8
04:47:08:07:de:17:23:13
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
  00:8e:6f:58:dd:68:bf:25:0a:e3:f9:17:70:d6:61:f3:53:a7:79:10 (sha1)
  71:6f:6a:76:17:9b:d6:2a:e7:5a:72:97:82:6d:26:86 (md5)
Distribution CRL:
  C=us, O=juniper, CN=CRL1
  http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: CRL signing, Certificate signing
Certificate identifier: entrust
Certificate version: 3
Serial number: 4355 925c
Issuer:
  Organization: juniper, Country: us
Subject:
  Organization: juniper, Country: us, Common name: First Officer
Validity:
  Not before: 2005 Oct 18th, 23:55:59 GMT
  Not after: 2008 Oct 19th, 00:25:59 GMT
Public key algorithm: rsaEncryption(1024 bits)
c0:a4:21:32:95:0a:cd:ec:12:03:d1:a2:89:71:8e:ce:4e:a6:f9:2f
1a:9a:13:8c:f6:a0:3d:c9:bd:9d:c2:a0:41:77:99:1b:1e:ed:5b:80
34:46:f8:5b:28:34:38:2e:91:7d:4e:ad:14:86:78:67:e7:02:1d:2e
19:11:b7:fa:0d:ba:64:20:e1:28:4e:3e:bb:6e:64:dc:cd:b1:b4:7a
ca:8f:47:dd:40:69:c2:35:95:ce:b8:85:56:d7:0f:2d:04:4d:5d:d8
42:e1:4f:6b:bf:38:c0:45:1e:9e:f0:b4:7f:74:6f:e9:70:fd:4a:78
da:eb:10:27:bd:46:34:33
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
  bc:78:87:9b:a7:91:13:20:71:db:ac:b5:56:71:42:ad:1a:b6:46:17 (sha1)
  23:79:40:c9:6d:a6:f0:ca:e0:13:30:d4:29:6f:86:79 (md5)
Distribution CRL:
  C=us, O=juniper, CN=CRL1
  http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: Key encipherment
Certificate identifier: entrust
Certificate version: 3
Serial number: 4355 925b

```

Issuer:
 Organization: juniper, Country: us
Subject:
 Organization: juniper, Country: us, Common name: First Officer
Validity:
 Not before: 2005 Oct 18th, 23:55:59 GMT
 Not after: 2008 Oct 19th, 00:25:59 GMT
Public key algorithm: rsaEncryption(1024 bits)
 ea:75:c4:f3:58:08:ea:65:5c:7e:b3:de:63:0a:cf:cf:ec:9a:82:e2
 d7:e8:b9:2f:bd:4b:cd:86:2f:f1:dd:d8:a2:95:af:ab:51:a5:49:4e
 00:10:c6:25:ff:b5:49:6a:99:64:74:69:e5:8c:23:5b:b4:70:62:8e
 e4:f9:a2:28:d4:54:e2:0b:1f:50:a2:92:cf:6c:8f:ae:10:d4:69:3c
 90:e2:1f:04:ea:ac:05:9b:3a:93:74:d0:59:24:e9:d2:9d:c2:ef:22
 b9:32:c7:2c:29:4f:91:cb:5a:26:fe:1d:c0:36:dc:f4:9c:8b:f5:26
 af:44:bf:53:aa:d4:5f:67
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
 46:71:15:34:f0:a6:41:76:65:81:33:4f:68:47:c4:df:78:b8:e3:3f (sha1)
 ee:cc:c7:f4:5d:ac:65:33:0a:55:db:59:72:2c:dd:16 (md5)
Distribution CRL:
 C=us, O=juniper, CN=CRL1
 http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: Digital signature

show security pki certificate-request

| | |
|---------------------------------|---|
| Syntax | show security pki certificate-request
<brief detail>
<certificate-id <i>certificate-id-name</i> > |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | Display information about manually generated local digital certificate requests that are stored in the router. |
| Options | <p>none—(same as brief) Display information about all local digital certificate requests.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>certificate-id <i>certificate-id-name</i>—(Optional) Display information about only the specified local digital certificate request</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear security pki certificate-request on page 1847 |
| List of Sample Output | show security pki certificate-request on page 1935
show security pki certificate-request detail on page 1935 |
| Output Fields | Table 68 on page 1934 lists the output fields for the show security pki certificate-request command. Output fields are listed in the approximate order in which they appear. |

Table 68: show security pki certificate-request Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------------|---|-------------------|
| Certificate identifier | Name of the digital certificate. | All levels |
| Certificate version | Revision number of the digital certificate. | detail |
| Issued to | Device that was issued the digital certificate. | none brief |
| Subject | <p>Details of the digital certificate holder organized using the distinguished name format. Possible subfields are:</p> <ul style="list-style-type: none"> Common name—Name of the authority. Organization—Organization of origin. Organizational unit—Department within an organization. State—State of origin. Country—Country of origin. | detail |
| Alternate subject | Domain name or IP address of the device related to the digital certificate. | detail |

Table 68: show security pki certificate-request Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|--|-----------------|
| Validity | Time period when the digital certificate is valid. Values are: <ul style="list-style-type: none"> • Not before—Time when the digital certificate becomes valid. • Not after—End time when the digital certificate becomes invalid. | All levels |
| Public key algorithm | Encryption algorithm used with the private key, such as rsaEncryption(1024 bits) . | All levels |
| Public key verification status | Public key verification status: Failed or Passed . The detail output also provides the verification hash. | All levels |
| Fingerprint | Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate. | detail |
| Use for key | Use of the public key, such as Certificate signing , CRL signing , Digital signature , or Key encipherment . | detail |

Sample Output

show security pki certificate-request

```

user@host> show security pki certificate-request
Certificate identifier: local-microsoft-2
Issued to: router2.juniper.net
Public key algorithm: rsaEncryption(1024 bits)
Public key verification status: Passed

```

show security pki certificate-request detail

```

user@host> show security pki certificate-request detail
Certificate identifier: local-entrust3
Certificate version: 3
Subject:
  Common name: router3.juniper.net
Alternate subject: router3.juniper.net
Public key algorithm: rsaEncryption(1024 bits)
Public key verification status: Passed
  fb:79:df:d4:a9:03:0f:d3:69:7e:c1:e4:27:35:9c:d9:b1:a2:47:78
  d2:6d:f3:e5:f4:68:4f:b3:04:45:88:57:99:82:39:a6:51:9e:5f:42
  23:3f:d7:6e:3d:a5:54:a9:b1:2d:6e:90:dd:12:8a:bf:ef:2b:20:50
  ba:f0:da:d9:0c:ad:5e:d6:c6:98:3a:ae:3f:90:dd:94:78:c1:ea:2e
  7c:f0:2d:d4:79:d4:cd:f0:52:df:5e:72:f2:e7:ae:66:f7:61:f4:bc
  72:57:3e:6c:6d:d3:24:58:8b:f4:ef:da:2a:6a:fa:eb:98:f8:34:84
  79:54:da:4f:d3:6f:52:1f
Fingerprint:
  7c:e8:f9:45:93:8d:a3:92:7f:18:29:02:f1:c8:e2:85:3d:ad:df:1f (sha1)
  00:4e:df:a0:6b:ad:8c:50:da:7c:a1:cf:5d:37:b0:ea (md5)
Use for key: Digital signature

```

show security pki crt

| | |
|---------------------------------|---|
| Syntax | show security pki crt
<brief detail>
<ca-profile <i>ca-profile-name</i> > |
| Release Information | Command introduced in Junos OS Release 8.1. |
| Description | Display information about the certificate revocation lists (CRLs) that are stored in the router. |
| Options | <p>none—(same as brief) Display information about all CRLs.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>ca-profile <i>ca-profile-name</i>—(Optional) Display CRL information about only the specified CA profile.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear security pki crt on page 1848 |
| List of Sample Output | show security pki crt on page 1937
show security pki crt detail on page 1937 |
| Output Fields | Table 69 on page 1936 shows the output fields for the show security pki crt command. Output fields are listed in the approximate order in which they appear. |

Table 69: show security pki crt Output Fields

| Field Name | Field Description | Level of Output |
|----------------|---|-----------------|
| CA profile | Name of the configured CA profile. | All levels |
| CRL version | Revision number of the certificate revocation list. | All levels |
| CRL number | Number of the certificate revocation list | All levels |
| CRL issuer | Device that was issued the certificate revocation list. | All levels |
| Issuer | <p>Details of the digital certificate holder organized using the distinguished name format. Possible subfields are:</p> <ul style="list-style-type: none"> • Common name—Name of the authority. • Organization—Organization of origin. • Organizational unit—Department within an organization. • State—State of origin. • Country—Country of origin. | detail |
| Effective date | Date and time the certificate revocation list becomes valid. | All levels |

Table 69: show security pki crl Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------|---|-----------------|
| Next update | Date and time the router will download the latest version of the certificate revocation list. | All levels |
| Revocation List | <p>List of digital certificates that have been revoked before their expiration date. Values are:</p> <ul style="list-style-type: none"> • Serial number—Unique serial number of the digital certificate • Revocation date—Date and time that the digital certificate was revoked. | detail |

Sample Output

show security pki crl

```
CA profile entrust
CRL version: V2
CRL number: 24
CRL issuer: C=CA, O=juniper
Effective date: 2006 May 31st, 05:35:25 GMT
Next update: 2006 Jun 1st, 06:35:25 GMT
```

show security pki crl detail

```
CA profile: entrust
CRL version: V2
CRL number: 24
Issuer:
  Organization: juniper, Country: ca
Validity:
  Effective date: 2006 May 31st, 05:35:25 GMT
  Next update: 2006 Jun 1st, 06:35:25 GMT
Revocation List:
  Serial number      Revocation date
  4451aca3 2006      May 25th, 09:13:38 GMT
  4451aca4 2006      May 25th, 10:11:33 GMT
  4451acb4 2006      May 29th, 11:28:54 GMT
  4451aceb 2006      May 29th, 11:29:01 GMT
  4451acfe 2006      May 29th, 11:29:17 GMT
  4451acff 2006      May 31st, 05:29:55 GMT
```

show security pki local-certificate

| | |
|---------------------------------|--|
| Syntax | show security pki local-certificate
<brief detail>
<certificate-id <i>certificate-id-name</i> >
<system-generated> |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | Display information about the local digital certificates and the corresponding public keys installed in the router. |
| Options | <p>none—(same as brief) Display information about all local digital certificates and corresponding public keys.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>certificate-id <i>certificate-id-name</i>—(Optional) Display information about only the specified the local digital certificate and corresponding public keys.</p> <p>system-generated—(Optional) Auto-generated self-signed certificate.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear security pki local-certificate on page 1850 |
| List of Sample Output | show security pki local-certificate on page 1939
show security pki local-certificate detail on page 1940 |
| Output Fields | Table 70 on page 1938 lists the output fields for the show security pki local-certificate command. Output fields are listed in the approximate order in which they appear. |

Table 70: show security pki local-certificate Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------------|--|-------------------|
| Certificate identifier | Name of the digital certificate. | All levels |
| Certificate version | Revision number of the digital certificate. | detail |
| Serial number | Unique serial number of the digital certificate. | detail |
| Issued by | Authority that issued the digital certificate. | none brief |
| Issued to | Device that was issued the digital certificate. | none brief |

Table 70: show security pki local-certificate Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------------|--|-----------------|
| Issuer | Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: <ul style="list-style-type: none"> • Common name—Name of the authority. • Organization—Organization of origin. • Organizational unit—Department within an organization. • State—State of origin. • Country—Country of origin. | detail |
| Subject | Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: <ul style="list-style-type: none"> • Common name—Name of the authority. • Organization—Organization of origin. • Organizational unit—Department within an organization. • State—State of origin. • Country—Country of origin. | detail |
| Alternate subject | Domain name or IP address of the device related to the digital certificate. | detail |
| Validity | Time period when the digital certificate is valid. Values are: <ul style="list-style-type: none"> • Not before—Start time when the digital certificate becomes valid. • Not after—End time when the digital certificate becomes invalid. | All levels |
| Public key algorithm | Encryption algorithm used with the private key, such as rsaEncryption (1024 bits) . | All levels |
| Public key verification status | Public key verification status: Failed or Passed . The detail output also provides the verification hash. | All levels |
| Signature algorithm | Encryption algorithm that the CA used to sign the digital certificate, such as sha1WithRSAEncryption . | detail |
| Fingerprint | Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate. | detail |
| Distribution CRL | Distinguished name information and URL for the certificate revocation list (CRL) server. | detail |
| Use for key | Use of the public key, such as Certificate signing , CRL signing , Digital signature , or Key encipherment . | detail |

Sample Output

show security pki local-certificate

```

user@host> show security pki local-certificate
Certificate identifier: local-entrust2
Issued to: router2.juniper.net, Issued by: juniper

```

```
Validity:
  Not before: 2005 Nov 21st, 23:28:22 GMT
  Not after: 2008 Nov 21st, 23:58:22 GMT
Public key algorithm: rsaEncryption(1024 bits)
Public key verification status: Passed
```

show security pki local-certificate detail

```
user@host> show security pki local-certificate detail
Certificate identifier: local-entrust3
Certificate version: 3
Serial number: 4355 94f9
Issuer:
  Organization: juniper, Country: us
Subject:
  Organization: juniper, Country: us, Common name: router3.juniper.net
Alternate subject: router3.juniper.net
Validity:
  Not before: 2005 Nov 21st, 23:33:58 GMT
  Not after: 2008 Nov 22nd, 00:03:58 GMT
Public key algorithm: rsaEncryption(1024 bits)
Public key verification status: Passed
fb:79:df:d4:a9:03:0f:d3:69:7e:c1:e4:27:35:9c:d9:b1:a2:47:78
d2:6d:f3:e5:f4:68:4f:b3:04:45:88:57:99:82:39:a6:51:9e:5f:42
23:3f:d7:6e:3d:a5:54:a9:b1:2d:6e:90:dd:12:8a:bf:ef:2b:20:50
ba:f0:da:d9:0c:ad:5e:d6:c6:98:3a:ae:3f:90:dd:94:78:c1:ea:2e
7c:f0:2d:d4:79:d4:cd:f0:52:df:5e:72:f2:e7:ae:66:f7:61:f4:bc
72:57:3e:6c:6d:d3:24:58:8b:f4:ef:da:2a:6a:fa:eb:98:f8:34:84
79:54:da:4f:d3:6f:52:1f
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
  61:3a:d0:b4:7a:16:9b:39:ba:81:3f:9d:ab:34:e5:c8:be:3b:a1:6d (sha1)
  60:a0:ff:58:05:4a:65:73:9d:74:3a:e1:83:6f:1b:c8 (md5)
Distribution CRL:
  C=us, O=juniper, CN=CRL1
  http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: Digital signature
```

show services cos statistics

| | |
|---------------------------------|---|
| Syntax | <pre>show services cos statistics <brief detail extensive> <diffserv forwarding-class> <interface <i>interface-name</i>> <service-set <i>service-set-name</i>> <summary></pre> |
| Release Information | Command introduced in Junos OS Release 8.1. |
| Description | Display the mapping of class-of-service (CoS) code point aliases to corresponding bit patterns and the mapping of forwarding class names to queue numbers as configured in CoS services for the AS PIC. |
| Options | <p>none—Display all services CoS statistics.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>diffserv forwarding-class—(Optional) Display only the selected information, either DiffServ codepoints or forwarding classes.</p> <p>interface <i>interface-name</i>—(Optional) Display statistics for the specified interface only.</p> <p>service-set <i>service-set-name</i>—(Optional) Display statistics for the specified service set only.</p> <p>summary—(Optional) Display summary of statistics on a per-interface basis.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services cos statistics on page 1942
show services cos statistics brief on page 1943
show services cos statistics detail on page 1943
show services cos statistics extensive on page 1943 |
| Output Fields | Table 71 on page 1941 describes the output fields for the show services cos statistics command. Output fields are listed in the approximate order in which they appear. |

Table 71: show services cos statistics Output Fields

| Field Name | Field Description | Level of Output |
|--------------------|----------------------------------|-----------------|
| Interface | Name of interface. | All levels |
| Service set | Name of service set. | All levels |
| DSCP | DiffServ code point bit pattern. | All levels |
| Packets in | Number of packets received. | All levels |

Table 71: show services cos statistics Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|--------------------------------|-----------------|
| Packets out | Number of packets transmitted. | All levels |
| Forwarding class | Forwarding class queue number. | All levels |

Sample Output

show services cos statistics

```

user@host> show services cos statistics
Interface: sp-1/0/0, Service set: scos
DSCP          Packets in      Packets out
000000          0             0
000001          0             0
000010          0             0
000011          0             0
000100          0             0
000101          0             0
000110          0             0
000111          0             0
001000          0             0
001001          0             0
001010          0             0
001011          0             0
001100          0             0
001101          0             0
001110          0             0
001111          0             0
010000          0             0
010001          0             0
010010          0             0
010011          0             0
010100          0             0
010101          0             0
010110          0             0
010111          0             0
011000          0             0
011001          0             0
011010          0             0
011011          0             0
011100          0             0
011101          0             0
011110          0             0
011111          0             0
100000          0             0
100001          0             0
100010          0             0
100011          0             0
100100          0             0
100101          0             0
100110          0             0
100111          0             0
101000          0             0
101001          0             0
101010          0             0

```


| | | |
|------------------|------------|-------------|
| 101011 | 0 | 0 |
| 101100 | 0 | 0 |
| 101101 | 0 | 0 |
| 101110 | 0 | 0 |
| 101111 | 0 | 0 |
| 110000 | 0 | 0 |
| 110001 | 0 | 0 |
| 110010 | 0 | 0 |
| 110011 | 0 | 0 |
| 110100 | 0 | 0 |
| 110101 | 0 | 0 |
| 110110 | 0 | 0 |
| 110111 | 0 | 0 |
| 111000 | 0 | 0 |
| 111001 | 0 | 0 |
| 111010 | 0 | 0 |
| 111011 | 0 | 0 |
| 111100 | 0 | 0 |
| 111101 | 0 | 0 |
| 111110 | 0 | 0 |
| 111111 | 0 | 0 |
| Forwarding class | Packets in | Packets out |
| 0 | 0 | 0 |
| 1 | 0 | 0 |
| 2 | 0 | 0 |
| 3 | 0 | 0 |
| 4 | 0 | 0 |
| 5 | 0 | 0 |
| 6 | 0 | 0 |
| 7 | 0 | 0 |
| 8 | 0 | 0 |
| 9 | 0 | 0 |
| 10 | 0 | 0 |
| 11 | 0 | 0 |
| 12 | 0 | 0 |
| 13 | 0 | 0 |
| 14 | 0 | 0 |
| 15 | 0 | 0 |

show services cos statistics brief

The output for the **show services cos statistics brief** command is identical to that for the **show services cos statistics** command.

show services cos statistics detail

The output for the **show services cos statistics detail** command is identical to that for the **show services cos statistics** command.

show services cos statistics extensive

The output for the **show services cos statistics extensive** command is identical to that for the **show services cos statistics** command.

show services crtp

| | |
|---------------------------------|---|
| Syntax | show services crtp
<extensive>
<interface <i>interface-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display Compressed Real-Time Transport Protocol (CRTP) extensive output. |
| Options | <p>none—Display CRTP extensive output for all interfaces.</p> <p>extensive—(Optional) Display extensive CRTP information.</p> <p>interface <i>interface-name</i>—(Optional) Display CRTP flow statistics for the specified interface. On M Series and T Series routers, a link services IQ (lsq-fpc/pic/port) or redundant link services IQ (rlsq-fpc/pic/port) interface.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services crtp extensive on page 1945 |
| Output Fields | Table 72 on page 1944 lists the output fields for the show services crtp command. Output fields are listed in the approximate order in which they appear. |

Table 72: show services crtp Output Fields

| Field Name | Field Description |
|---------------------------------|--|
| Interface | Name of the physical interface. |
| Port minimum
Port maximum | Compression is applied to UDP packets with even ports in the specified range. |
| Maximum UDP compressed sessions | Maximum value of a context identifier in the space of context identifiers allocated for UDP. |
| CRTP maximum period | Maximum interval between full headers. Suggested value is 256. |
| CRTP maximum time | Maximum time interval between full headers. Suggested value is 5 seconds. |
| Compression ratio | Ratio of received packet size to compressed packet size, in percentage. For example, if the packet size is 100 bytes when it is received, and is 40 bytes after compression, the compression ratio is $100 \div 40 / 100 * 100 = 60\%$. |
| Decompression ratio | Ratio of received packet size to decompressed packet size, in percentage. For example, if the packet size is 40 bytes when it is received, and is 100 bytes after compression, the decompression ratio is $100 \div 40 / 100 * 100 = 60\%$. |

Table 72: show services crtp Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------------|--|
| Discards | Number of frames that the incoming packet match code discarded because they were not recognized. |
| Sessions | Total number of active CRTP sessions. |
| IP bytes | Number of IP bytes sent and received. |
| Compressed bytes | Number of compressed IP header bytes sent and received. |
| CRTP packets | Number of CRTP packets sent and received. |
| CUDP/CNTCP packets | Number of compressed UDP packets and compressed non-TCP packets sent and received. |
| Full header packets | Number of full header packets sent and received. Full header packets communicate the uncompressed IP header plus any following headers and data to establish the uncompressed header state in the decompressor for a particular context. |
| Context state packet | Number of context state packets sent and received. Context state packets are sent from the decompressor to the compressor to communicate a list of context IDs for which synchronization is lost or might be lost. |
| IP packets | Number of IP packets sent and received. |
| Compressed packets | Number of compressed packets sent and received. |

Sample Output

show services crtp extensive

```

user@host> show services crtp extensive
Interface: lsq-1/1/0.1
  Port minimum: 2000, Port maximum: 64009
  Maximum UDP compressed sessions: 256
  CRTP maximum period: 256, CRTP maximum time: 5
  Compression ratio: 0, Decompression ratio: 0, Discards: 0
  CRTP stats
    Receive      Transmit
  Sessions           1           1
  IP bytes           60           60
  Compressed bytes   61           60
  CRTP packets       0           0
  CUDP/CNTCP packets 0           0
  Full header packets 1           1
  Context state packets 0           0
  IP packets         1           1
  Compressed packets 1           1

```

show services crtp flows

| | |
|---------------------------------|---|
| Syntax | show services crtp flows
<interface <i>interface-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display Compressed Real-Time Transport Protocol (CRTP) flows. |
| Options | <p>none—Display CRTP flows for all interfaces.</p> <p>interface <i>interface-name</i>—(Optional) Display CRTP flows for the specified interface. On M Series and T Series routers, a link services IQ (lsq-<i>fpc/pic/port</i>) or redundant link services IQ (rlsq-<i>fpc/pic/port</i>) interface.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services crtp flows on page 1946 |
| Output Fields | Table 73 on page 1946 lists the output fields for the show services crtp flows command. Output fields are listed in the approximate order in which they appear. |

Table 73: show services crtp flows Output Fields

| Field Name | Field Description |
|--------------------|---|
| Interface | Name of the physical interface. |
| Flow | Received or transmitted flow. |
| Source | IP source address. |
| Destination | IP destination address. |
| SSRC ID | Synchronization source (SSRC) identifier. One of the fields in the RTP header used to select the context. The SSRC identifier is a randomly chosen value unique within a particular CRTP session. |
| Ctx ID | Session context ID. Indicates the session context in which to interpret the packet. The decompressor can use the context ID to index its table of stored session contexts directly. |

Sample Output

show services crtp flows

```

user@host> show services crtp flows
Interface: lsq-1/1/0.1
  Flow      Source           Destination      SSRC ID  Ctx ID
  Receive   60.1.1.3:28004      80.1.1.3:26000   123      0
  Transmit  80.1.1.3:26000      60.1.1.3:28004   123      2

```


show services ids

Syntax show services ids (destination-table | pair-table | source-table)
<brief | extensive | terse>
<destination-prefix *destination-prefix-name*>
<interface *interface-name*>
<limit *number*>
<order (anomalies | bytes | flows | packets)>
<service-set *service-set-name*>
<source-prefix *source-prefix-name*>
<threshold *number*>

Release Information Command introduced before Junos OS Release 7.4.

Description Display information about intrusion detection service (IDS) events. All events gathered by IDS are reported as anomalies. For example, events such as **create forward or watch flow**, **FTP passive**, and **FTP active** are genuinely allowed by the stateful firewall but are logged as anomalies to track the rates and number for these events.

Options **destination-table**—Display information for an address under possible attack.

pair-table—Display information for a particular suspected attack source and destination address pair.

source-table—Display information for an address that is a suspected attacker.

brief | extensive | terse—(Optional) Display the specified level of output.

destination-prefix *destination-prefix-name*—(Optional) Display information for a particular destination prefix.

interface *interface-name*—(Optional) On M Series and T Series routers, the *interface-name* can be *sp-fpc/pic/port* or *rspnumber*.

limit *number*—(Optional) Maximum number of entries to display. By default, all tables display the top 32 entries sorted by the number of events for the criteria chosen. To display additional entries, configure the limit option to set up to 256 entries.

order—(Optional) Display events according to one of the following table-ordering criteria. The default is anomalies.

- **anomalies**—Display information for particular anomalies.
- **bytes**—Order output by number of bytes received.
- **flows**—Order output by number of flows.
- **packets**—Order output by number of packets received.

service-set *service-set-name*—(Optional) Display information about a particular service set.

source-prefix *source-prefix-name*—(Optional) Display information about a particular source prefix.

threshold *number*—(Optional) Limit the display to events with this number of anomalies, bytes, flows, or packets, whichever criterion you specify for order. For example, to display all events with more than 100 flows, specify order flows and threshold 100.

Required Privilege Level view

List of Sample Output [show services ids destination-table on page 1952](#)
[show services ids destination-table extensive on page 1952](#)
[show services ids destination-table extensive order anomalies on page 1952](#)
[show services ids pair-table extensive on page 1953](#)
[show services ids pair-table extensive limit on page 1953](#)
[show services ids source-table extensive on page 1954](#)
[show services ids source-table extensive limit on page 1954](#)

Output Fields [Table 74 on page 1949](#) lists the output fields for the **show services ids** command. Output fields are listed in the approximate order in which they appear.

Table 74: show services ids Output Fields

| Field Name | Field Description | Output Level |
|-----------------------|--|--------------|
| Interface | Name of an adaptive services interface. | All levels |
| Service set | Name of a service set. Individual empty service sets are not displayed, but if no service set has any flows, a flow table header is printed for each service set. | All levels |
| Sorting order | Primary mode to display information: Anomalies , Bytes , Flows , or Packets . | All levels |
| Source address | Name of the source address. | All levels |
| Dest address | Name of the destination address. | All levels |
| Time | Total time the information has been in the table. | All levels |
| Flags | Flags can be Forced , F (terse output only), SYNcookie , S (terse output only), Forced+SYNcookie , and F+S (terse output only). The SYNcookie flag is visible only in the destination table. | All levels |
| Application | Configured application, such as FTP or Telnet . | All levels |
| Bytes | Total number of bytes sent from the source to the destination address, in thousands (k) or millions (m). | All levels |
| Packets | Total number of packets sent from the source to the destination address, in thousands (k) or millions (m). | All levels |
| Flows | Total number of flows of packets sent from the source to the destination address, in thousands (k) or millions (m). | All levels |

Table 74: show services ids Output Fields (*continued*)

| Field Name | Field Description | Output Level |
|----------------------------|---|------------------|
| Anomalies | Total number of packets in the anomaly table, in thousands (k) or millions (m). | All levels |
| Anomaly description | <p>One or more of the following types of anomalies. For more information, see the detailed descriptions in the stateful firewall section of the <i>Junos OS System Log Messages Reference</i>.</p> <ul style="list-style-type: none"> • First packet of TCP session not SYN • ICMP echo request dropped, because sequence number duplicated • ICMP echo reply dropped. No matching sequence number • ICMP echo request dropped. Too many echo requests without echo reply • ICMP header length check failed • ICMP packet length greater than 64K • IP fragment assembly timeout • IP fragment length error • IP fragment overlap • IP packet length greater than 64K • IP packet too short • IP packet with broadcast destination address • IP packet with checksum error • IP packet with incorrect length • IP packet with TTL equal to 0 | extensive |

Table 74: show services ids Output Fields (*continued*)

| Field Name | Field Description | Output Level |
|------------------------------------|--|--------------|
| Anomaly description
(continued) | <ul style="list-style-type: none"> • IP packet with version other than 4 • Land attack (IP src address = dest address) • No matching SFW rule; attempting to create discard flow • Number of open sessions exceeds IDS limit; packet dropped • Packet rate exceeds IDS limit; packet dropped • Session creation rate exceeds IDS limit; packet dropped • SFW application message too long • SFW discard packet contains non-configured IP option types • SFW drop packet because of discard flow • SFW dropped TCP watch packet • SFW rules request FTP active mode data packets to be accepted; attempting to create forward flow • SFW rules request FTP passive mode data packets to be accepted; attempting to create forward flow • SFW rules request packet to be accepted; attempting to create forward or watch flow • SFW rules request packet to be discarded; attempting to create discard flow • SFW rules request packet to be rejected; attempting to create reject flow • SFW discard flow requires packet to be dropped • SFW SYN defense • Smurf attack (ping to IP broadcast address) • TCP FIN/RST or SYN/(URG FIN RST) flags set • TCP header length check failed • TCP port scan (port not in LISTEN state) • TCP seq number zero and FIN/PSH/RST flags set • TCP seq number zero and no flags set • TCP source or destination port zero • TCP SYN flood attack • UDP header length check failed • UDP port scan (port not in LISTEN state) • UDP source or destination port zero | extensive |
| Count | Number of times that a particular anomaly occurred, in thousands (k) or millions (M). | extensive |
| Rate (eps) | Anomaly events per second. The IDS subsystem attempts to maintain a weighted average of rates, which might not reflect the exact incoming rate of attack at low rates. However, at high rates exceeding 160 events per second, the rates generally match. | extensive |
| Elapsed | Time since the same type of event last occurred. | extensive |
| Total IDS table entries | Number of entries in the IDS table. This number is not necessarily the sum of all entries displayed. | All levels |

Table 74: show services ids Output Fields (*continued*)

| Field Name | Field Description | Output Level |
|--|---|--------------|
| Total failed IDS table entry insertions | Number of IDS entries not allowed into the table because the table was full | All levels |
| Total number of events (closed flows and anomalies detected) | Total number of events since the system was started or since the show ids services command was executed. | All levels |

Sample Output

show services ids destination-table

```

user@host> show services ids destination-table
Interface: sp-1/3/0, Service set: null-sfw
Sorting order: Packets
Source address      Dest address  Time    Flags           Application

any                ->  10.58.255.146  36m12s SYN cookie
Bytes:  35.0 m, Packets:  822.0 k, Flows:  274.0 k, Anomalies: 2251.0 k

Total IDS table entries: 87
Total failed IDS table entry insertions 0
Total number of events (closed flows and anomalies detected): 2606018

```

show services ids destination-table extensive

```

user@host> show services ids destination-table extensive
Interface: sp-1/3/0, Service set: null-sfw
Sorting order: Packets
Source address      Dest address  Time    Flags           Application

any                ->  10.58.255.146  35m52s SYN cookie
Bytes:  34.0 m, Packets:  798.0 k, Flows:  266.0 k, Anomalies: 2251.0 k
Anomalies
First packet of TCP session not SYN          Count Rate(eps) Elapsed
TCP source or destination port zero          634.0 k  154.6      3m37s
UDP source or destination port zero          633.0 k  170.0      3m37s
ICMP header length check failed              2875    0.9        3m37s
IP fragment assembly timeout                 820.0 k  12.8       3m18s
UDP header length check failed                385     0.5        3m53s
TCP header length check failed                383     0.5        3m53s

Total IDS table entries:
87
Total failed IDS table entry insertions
0
Total number of events (closed flows and anomalies detected):
2598063

```

show services ids destination-table extensive order anomalies

```

user@host> show services ids destination-table extensive order anomalies

```

```

Interface: sp-0/2/0, Service set: ssl
IDS sorting order: Anomalies
Source address      Dest address      Time Flags      Application
15.1.1.1            -> 15.99.1.1        1m28s          junos-ftp
Bytes: 1065, Packets: 18, Flows: 1, Anomalies: 10
Anomaly description
creating forward or watch flow          Count  Rate(eps)  Elapsed
Number of open sessions exceeds IDS limit 9      0.8        18s

Total IDS table entries:                3
Total failed IDS table entry insertions 0
Total number of events (closed flows and anomalies): 11

```

show services ids pair-table extensive

```

user@host> show services ids pair-table extensive
Interface: sp-3/2/0, Service set: ss_all_limits
IDS sorting order: Packets
Source address      Dest address      Time Flags      Application
15.1.1.4            -> 15.99.1.4        2m20s          junos-ftp

Bytes: 5.7k, Packets: 102.0, Flows: 41.0, Anomalies: 462.0
Anomaly description          Count  Rate  Elapsed
creating forward or watch flow 41.0   8.8   2m17s

Packet rate exceeds IDS src limit 21.0   7.1   2m17s

Session creation rate exceeds IDS src limit 359.0  99.7   2m16s

TCP SYN flood attack          41.0   1.9   1m30s

Total IDS table entries:                3
Total failed IDS table entry insertions 0
Total number of events (closed flows and anomalies): 462

```

show services ids pair-table extensive limit

```

user@host> show services ids pair-table extensive limit 3
Interface: sp-1/3/0, Service set: null-sfw
Sorting order: Packets
Source address      Dest address      Time  Flags      Application
10.58.255.18        -> 10.58.255.146    38m41s SYN cookie
Bytes: 286.0 m, Packets: 2823.0 k, Flows: 324.0 k, Anomalies: 387.0 k
Anomalies          Count  Rate(eps)  Elapsed
First packet of TCP session not SYN 160.0 k 0.1        25s
TCP source or destination port zero 69.0 k 14.1       6m26s
UDP source or destination port zero 68.0 k 12.7       6m26s
ICMP header length check failed 318    0.1        7m6s
IP fragment assembly timeout 88.0 k 1.3        6m7s
UDP header length check failed 39     0.0        6m58s
TCP header length check failed 46     0.0        6m45s

10.58.255.23        -> 10.58.255.146    18m48s SYN cookie
Bytes: 104.0 m, Packets: 421.0 k, Flows: 230, Anomalies: 124.0 k
Anomalies          Count  Rate(eps)  Elapsed
TCP source or destination port zero 37.0 k 9.8        6m26s
UDP source or destination port zero 37.0 k 8.4        6m26s
IP fragment assembly timeout 48.0 k 1.0        6m7s
ICMP header length check failed 190    0.2        6m47s

```

```

UDP header length check failed          29    0.0    6m51s
TCP header length check failed          23    0.0    6m59s

10.58.255.25  ->  10.58.255.146  18m48s SYN cookie
Bytes: 104.0 m, Packets: 420.0 k, Flows: 232, Anomalies: 123.0 k
Anomalies
TCP source or destination port zero      37.0 k    9.8    6m26s
UDP source or destination port zero      37.0 k    8.6    6m26s
IP fragment assembly timeout            48.0 k    1.5     6m7s
ICMP header length check failed          173     0.1    6m43s
UDP header length check failed           24     0.0    6m43s
TCP header length check failed           19     0.0    6m56s

Total IDS table entries:
87
Total failed IDS table entry insertions
0
Total number of events (closed flows and anomalies detected):
2659291

```

show services ids source-table extensive

```

user@host> show services ids source-table extensive
Interface: sp-3/2/0, Service set: ss_all_limits
IDS sorting order: Packets
Source address      Dest address      Time Flags      Application
15.1.1.4            ->               any      2m43s         junos-ftp

Bytes: 5.7k, Packets: 102.0, Flows: 41.0, Anomalies: 462.0
Anomaly description      Count    Rate    Elapsed
creating forward or watch flow      41.0      8.8    2m40s

Packet rate exceeds IDS src limit      21.0      7.1    2m40s

Session creation rate exceeds IDS src limit      359.0    99.7    2m39s

TCP SYN flood attack      41.0      1.9    1m53s

Total IDS table entries:          3
Total failed IDS table entry insertions      0
Total number of events (closed flows and anomalies):      462

```

show services ids source-table extensive limit

```

user@host> show services ids source-table extensive limit 3
Interface: sp-1/3/0, Service set: null-sfw
Sorting order: Packets
Source address      Dest address      Time  Flags      Application
10.58.255.18        ->               any   40m 0s SYN cookie
Bytes: 250.0 m, Packets: 1978.0 k, Flows: 356.0 k, Anomalies: 387.0 k
Anomalies
TCP source or destination port zero      37.0 k    9.8    6m26s
First packet of TCP session not SYN      160.0 k    0.0     40s
TCP source or destination port zero      69.0 k   62.5    7m45s
UDP source or destination port zero      68.0 k   56.2    7m45s
ICMP header length check failed          319     0.1    7m49s
IP fragment assembly timeout            89.0 k    4.4    7m26s
UDP header length check failed           39     0.0    8m17s

```

```

TCP header length check failed                46      0.0      8m4s

10.58.255.30  ->                any  20m 7s SYN cookie
Bytes: 107.0 m, Packets: 427.0 k, Flows: 264, Anomalies: 125.0 k
Anomalies
UDP source or destination port zero          38.0 k    65.5    7m45s
TCP source or destination port zero          37.0 k    38.1    7m45s
IP fragment assembly timeout                 49.0 k     4.1    7m26s
TCP header length check failed                24      0.0    9m23s
ICMP header length check failed              165      0.1     8m6s
UDP header length check failed                26      0.0    8m13s

10.58.255.17  ->                any  20m10s SYN cookie
Bytes: 107.0 m, Packets: 426.0 k, Flows: 262, Anomalies: 125.0 k
Anomalies
TCP source or destination port zero          38.0 k    55.     7m45s
UDP source or destination port zero          38.0 k    55.1    7m45s
ICMP header length check failed              147      0.1    7m50s
IP fragment assembly timeout                 49.0 k     2.8    7m26s
TCP header length check failed                22      0.0    9m33s
UDP header length check failed                22      0.0     8m1s

Total IDS table entries:
87
Total failed IDS table entry insertions
0
Total number of events (closed flows and anomalies detected):
2691423
Interface: sp-1/3/0, Service set: blue
NAT pool      Address      Port      Ports in use
d2-pool      10.59.16.100-10.59.16.100  4000-4002      1

```

show services inline nat pool

| | |
|---------------------------------|---|
| Syntax | <code>show services inline nat pool</code>
<code><pool <i>pool--name</i>></code> |
| Release Information | Command introduced in Junos OS Release 11.4. |
| Description | Display information about inline Network Address Translation (NAT) pool. |
| Options | <i>pool-name</i> —Display information about the specified services-inline interface NAT pool. |
| Required Privilege Level | view |
| List of Sample Output | show services inline nat pool on page 1956 |
| Output Fields | Table 75 on page 1956 lists the output fields for the <code>show services inline nat pool</code> command. Output fields are listed in the order in which they appear. |

Table 75: show services inline nat pool Output Fields

| Field Name | Field Description |
|-------------------------|--|
| Interface | Name of an <code>si</code> interface hosted on a Trio-based line card. |
| NAT pool | Name of the pool used for address translations. |
| Translation type | Translation type specified in the applicable NAT rule for the service set. |
| Address range | Starting and ending public NAT addresses available for translation. |
| NATed packets | Number of packets translated for the specified pool. |
| un-NATed packets | Number of received packets that were not translated. |
| Errors | Number of packets with translation errors. |

Sample Output

show services inline nat pool

```

user@host> show services inline nat pool p1
Interface: si-5/0/0, Service set: ss-inat
NAT pool: p1, Translation type: BASIC NAT44
Address range: 20.20.20.0-20.20.20.255
NATed packets: 0, Un-NATed packets: 0, Errors: 0

```

show services inline nat statistics

| | |
|---------------------------------|---|
| Syntax | <code>show services inline nat statistics</code>
<code><interface <i>interface-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 11.4. |
| Description | Display information about inline Network Address Translation (NAT) address translations. |
| Options | <i>interface-name</i> —(Optional) Display information about the specified NAT services-inline interface only. When a specific interface is not specified, statistics for all services-inline interfaces are shown. |
| Required Privilege Level | view |
| List of Sample Output | show services inline nat statistics on page 1957 |
| Output Fields | Table 76 on page 1957 lists the output fields for the show services inline nat statistics command. Output fields are listed in the order in which they appear. |

Table 76: show services inline nat statistics Output Fields

| Field Name | Field Description | Level of Output |
|-----------------------------------|--|-----------------|
| Service PIC | Name of an si interface hosted on a Trio-based line card. | All levels |
| Slow path packets received | Number of ICMP exception packets received for NAT translation. | All levels |
| Slow path packets dropped | Number of received ICMP exception packets that were dropped. | All levels |

Sample Output

show services inline nat statistics

```

user@host> show services inline nat statistics
Service PIC Name                               :si-5/0/0

Slow path packets received                     :0
Slow path packets dropped                      :0

```

show services ipsec-vpn certificates

| | |
|---------------------------------|--|
| Syntax | <code>show services ipsec-vpn certificates</code>
<code><brief detail></code>
<code><service-set service-set></code> |
| Release Information | Command introduced in Junos OS Release 7.5. |
| Description | (Adaptive services interfaces only) Display local and remote certificates installed in the IPsec configuration memory cache that are used for the IKE negotiation. |
| Options | <p>none—(same as brief) Display information about local and remote certificates associated with all service sets.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>service-set service-set—(Optional) Display information about local and remote certificates associated with only the specified service set.</p> |
| Required Privilege Level | view |
| List of Sample Output | show security ipsec-vpn certificates on page 1959
show security ipsec-vpn certificates detail on page 1959 |
| Output Fields | Table 77 on page 1958 lists the output fields for the show services ipsec-vpn certificates command. Output fields are listed in the approximate order in which they appear. |

Table 77: show services ipsec-vpn certificates Output Fields

| Field Name | Field Description | Level of Output |
|--------------------------------|---|-------------------|
| Service set | Name of the IPsec service set. | All levels |
| Total entries | Number of certificate cache entries. | All levels |
| Certificate cache entry | Identification number of the certificate cache entry. | All levels |
| Flags | Information about the digital certificate, including whether the certificate is a root certificate and trusted. | none brief |
| Issued to | Device that was issued the digital certificate. | none brief |
| Issued by | Authority that issued the digital certificate. | none brief |
| Certificate version | Revision number of the digital certificate. | detail |
| Serial number | Unique serial number of the digital certificate. | detail |
| Alternate subject | Domain name or IP address of the device related to the digital certificate. | All levels |

Table 77: show services ipsec-vpn certificates Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------|--|-------------------|
| Validity | Time period when the digital certificate is valid. Values are: <ul style="list-style-type: none"> • Not before—Start time when the digital certificate becomes valid. • Not after—End time when the digital certificate becomes invalid. | none brief |
| Public key algorithm | Specifies the encryption algorithm used with the private key, such as rsaEncryption (1024 bits) . | detail |
| Signature algorithm | Encryption algorithm that the CA used to sign the digital certificate, such as sha1WithRSAEncryption . | detail |
| Fingerprint | Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate. | detail |
| Distribution CRL | Distinguished name information and the URL for the certificate revocation list (CRL) server. | detail |
| Use for key | Use of the public key, such as Certificate signing, CRL signing, Digital signature, or Key encipherment . | detail |

Sample Output

show security ipsec-vpn certificates

```

user@host> show services ipsec-vpn certificates
Service set: serviceset-dynamic-BiEspsha3des, Total entries: 3
Certificate cache entry: 3
  Flags: Non-root Trusted
  Issued to: router3.juniper.net, Issued by: juniper
  Alternate subject: router3.juniper.net
  Validity:
    Not before: 2005 Nov 21st, 23:33:58 GMT
    Not after: 2008 Nov 22nd, 00:03:58 GMT

Certificate cache entry: 2
  Flags: Non-root Trusted
  Issued to: router2.juniper.net, Issued by: juniper
  Alternate subject: router2.juniper.net
  Validity:
    Not before: 2005 Nov 21st, 23:28:22 GMT
    Not after: 2008 Nov 21st, 23:58:22 GMT

Certificate cache entry: 1
  Flags: Root Trusted
  Issued to: juniper, Issued by: juniper
  Validity:
    Not before: 2005 Oct 18th, 23:54:22 GMT
    Not after: 2025 Oct 19th, 00:24:22 GMT

```

show security ipsec-vpn certificates detail

```

user@host> show services ipsec-vpn certificates detail

```

Service set: serviceset-dynamic-BiEspsha3des, Total entries: 3

Certificate cache entry: 3
Certificate version: 3
Serial number: 4355 94f9
Alternate subject: router3.juniper.net
Public key algorithm: rsaEncryption
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
61:3a:d0:b4:7a:16:9b:39:ba:81:3f:9d:ab:34:e5:c8:be:3b:a1:6d (sha1)
60:a0:ff:58:05:4a:65:73:9d:74:3a:e1:83:6f:1b:c8 (md5)
Distribution CRL:
C=us, O=juniper, CN=CRL1
http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: Digital signature

Certificate cache entry: 2
Certificate version: 3
Serial number: 4355 94f8
Alternate subject: router2.juniper.net
Public key algorithm: rsaEncryption
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
30:c3:a4:04:da:33:9d:60:23:5a:48:75:48:2c:f0:c6:96:6c:31:fa (sha1)
9a:a2:ce:ef:7e:10:80:a0:c8:4d:2f:e7:e1:d3:69:9d (md5)
Distribution CRL:
C=us, O=juniper, CN=CRL1
http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: Digital signature

Certificate cache entry: 1
Certificate version: 3
Flags: Root
Serial number: 4355 9235
Public key algorithm: rsaEncryption
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
00:8e:6f:58:dd:68:bf:25:0a:e3:f9:17:70:d6:61:f3:53:a7:79:10 (sha1)
71:6f:6a:76:17:9b:d6:2a:e7:5a:72:97:82:6d:26:86 (md5)
Distribution CRL:
C=us, O=juniper, CN=CRL1
http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: CRL signing, Certificate signing

show services ipsec-vpn ike security-associations

| | |
|---------------------------------|--|
| Syntax | show services ipsec-vpn ike security-associations
<brief detail>
<peer-address> |
| Release Information | Command introduced before Junos OS Release 7.4.
Statistics for Internet Key Exchange (IKE) security associations for each services PIC introduced in Junos OS Release 12.1. |
| Description | (Adaptive services interface only) Display information for Internet Key Exchange (IKE) security associations. If no security association is specified, the information for all security associations is displayed. |
| Options | none —(same as brief) Display standard information for all IPsec security associations.

brief detail —(Optional) Display the specified level of output.

peer-address —(Optional) Display information about a particular security association address. |
| Required Privilege Level | view |
| List of Sample Output | show services ipsec-vpn ike security-associations on page 1963
show services ipsec-vpn ike security-associations detail on page 1964 |
| Output Fields | Table 78 on page 1961 lists the output fields for the show services ipsec-vpn ike security-associations command. Output fields are listed in the approximate order in which they appear. |

Table 78: show services ipsec-vpn ike security-associations Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------|---|-----------------|
| IKE peer | Remote end of the IKE negotiation. | detail |
| Role | Part played in the IKE session. The router triggering the IKE negotiation is the initiator, and the router accepting the first IKE exchange packets is the responder. | detail |
| Remote Address | Responder's address. | none specified |
| State | State of the IKE security association: <ul style="list-style-type: none"> • Matured—IKE security association is established. • Not matured—The IKE security association is in the process of negotiation. | none specified |
| Initiator cookie | When the IKE negotiation is triggered, a random number is sent to the remote node. | All levels |

Table 78: show services ipsec-vpn ike security-associations Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------------|---|-----------------|
| Responder cookie | <p>The remote node generates its own random number and sends it back to the initiator as a verification that the packets were received.</p> <p>Of the numerous security services available, protection against denial of service (DoS) is one of the most difficult to address. A “cookie” or anticlogging token (ACT) is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity. An exchange prior to CPU-intensive public key operations can thwart some DoS attempts (such as simple flooding with invalid IP source addresses).</p> | All levels |
| Exchange type | <p>Specifies the number of messages in an IKE exchange, and the payload types that are contained in each message. Each exchange type provides a particular set of security services, such as anonymity of the participants, perfect forward secrecy of the keying material, and authentication of the participants. Junos OS supports two types of exchanges:</p> <ul style="list-style-type: none"> • Main—The exchange is done with six messages. Main encrypts the payload, protecting the identity of the neighbor. • Aggressive—The exchange is done with three messages. Aggressive does not encrypt the payload, leaving the identity of the neighbor unprotected. • IKEv2—The exchange is negotiated using IKE version 2. | All levels |
| PIC | The services PIC for which the IKE security associations are displayed. | All levels |
| Authentication method | Type of authentication determines which payloads are exchanged and when they are exchanged. The Junos OS supports only pre-shared keys . | detail |
| Local | Prefix and port number of the local end. | detail |
| Remote | Prefix and port number of the remote end. | detail |
| Lifetime | Number of seconds remaining until the IKE security association expires. | detail |
| Algorithms | <p>Header for the IKE algorithms output.</p> <ul style="list-style-type: none"> • Authentication—(detail output only) Type of authentication algorithm used: md5 or sha1 • Encryption—(detail output only) Type of encryption algorithm used: des-cbc, 3des-cbc, or None. • Pseudo random function—Function that generates highly unpredictable random numbers: hmac-md5 or hmac-sha1. | detail |
| Traffic statistics | <p>Number of bytes and packets received and transmitted on the IKE security association.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the IKE security association. • Input packets, Output packets—Number of packets received and transmitted on the IKE security association. | detail |

Table 78: show services ipsec-vpn ike security-associations Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|--|-----------------|
| Flags | Notification to the key management process of the status of the IKE negotiation: <ul style="list-style-type: none"> caller notification sent—Caller program notified about the completion of the IKE negotiation. waiting for done—Negotiation is done. The library is waiting for the remote end retransmission timers to expire. waiting for remove—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation. waiting for policy manager—Negotiation is waiting for a response from the policy manager. | detail |
| IPsec security associates | Number of IPsec security associations created and deleted with this IKE security association. | detail |
| Phase 2 negotiations in progress | Number of phase 2 negotiations in progress and status information: <ul style="list-style-type: none"> Negotiation type—Type of phase 2 negotiation. The Junos OS currently supports quick mode. Message ID—Unique identifier for a phase 2 negotiation. Local identity—Identity of the local phase 2 negotiation. The format is <i>id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)</i>. Remote identity—Identity of the remote phase 2 negotiation. The format is <i>id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)</i>. Flags—Notification to the key management process of the status of the IKE negotiation: <ul style="list-style-type: none"> caller notification sent—Caller program notified about the completion of the IKE negotiation. waiting for done—Negotiation is done. The library is waiting for the remote end retransmission timers to expire. waiting for remove—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation. waiting for policy manager—Negotiation is waiting for a response from the policy manager. | detail |

Sample Output

show services ipsec-vpn ike security-associations

```

user@host> show services ipsec-vpn ike security-associations
Remote Address  State      Initiator cookie  Responder cookie  Exchange type
-----
6.6.6.1         Matured    062d291d21275fc7  82ef00e3d1f1c981  Main
6.6.6.2         Matured    cd6d581d7bb1664d  88a707779f3ad8d1  Main
6.6.6.3         Matured    86621051e3e78360  6bc5cc83fd67baa4  IKEv2
PIC: sp-0/3/0
6.6.6.7         Matured    565e2813075e6fdb  67886757a74edcd6  IKEv2

```

show services ipsec-vpn ike security-associations detail

```
user@host> show services ipsec-vpn ike security-associations detail
IKE peer 3.1.0.2
  Role: Responder, State: Matured
  Initiator cookie: d91c9f20f78e1d4e, Responder cookie: 727a04ed8d5021a1
  Exchange type: IKEv2, Authentication method: Pre-shared-keys
  Local: 4.1.0.2:500, Remote: 3.1.0.2:500
  Lifetime: Expires in 1357 seconds
  Algorithms:
    Authentication      : sha1
    Encryption          : 3des-cbc
    Pseudo random function: hmac-sha1
  Traffic statistics:
    Input bytes  :          22244
    Output bytes :          22236
    Input packets:           263
    Output packets:          263
  Flags: Caller notification sent
  IPsec security associations: 0 created, 0 deleted
  Phase 2 negotiations in progress: 0

IKE peer 4.4.4.4
  Role: Initiator, State: Matured
  Initiator cookie: cf22bd81a7000001, Responder cookie: fe83795c2800002e
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 4.4.4.5:500, Remote: 4.4.4.4:500
  Lifetime: Expires in 187 seconds
  Algorithms:
    Authentication      : md5
    Encryption          : 3des-cbc
    Pseudo random function: hmac-md5
  Traffic statistics:
    Input bytes  :          1000
    Output bytes :          1280
    Input packets:           5
    Output packets:           9
  Flags: Caller notification sent
  IPsec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

Negotiation type: Quick mode, Role: Initiator, Message ID: 3582889153
  Local: 4.4.4.5:500, Remote: 4.4.4.4:500
  Local identity: ipv4_subnet(tcp:80,[0..7]=10.1.1.0/24)
  Remote identity: ipv4_subnet(tcp:100,[0..7]=10.1.2.0/24)
  Flags: Caller notification sent, Waiting for done
```

show services ipsec-vpn ipsec security-associations

| | |
|---------------------------------|---|
| Syntax | show services ipsec-vpn ipsec security-associations
<brief detail extensive>
<service-set <i>service-set-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (Adaptive services interface only) Display IPsec security associations for the specified service set. If no service set is specified, the security associations for all service sets are displayed. |
| Options | <p>none—Display standard information about IPsec security associations for all service sets.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>service-set <i>service-set-name</i>—(Optional) Display information about a particular service set.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services ipsec-vpn ipsec security associations extensive on page 1968 |
| Output Fields | Table 79 on page 1965 lists the output fields for the show services ipsec-vpn ipsec security-associations command. Output fields are listed in the approximate order in which they appear. |

Table 79: show services ipsec-vpn ipsec security-associations Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------------|---|-------------------------|
| Service set | Name of the service set for which the IPsec security associations are defined. If appropriate, includes the outside service interface VRF name. | All levels |
| Rule | Name of the rule set applied to the security association. | detail extensive |
| Term | Name of the IPsec term applied to the security association. | detail extensive |
| Tunnel index | Numeric identifier of the specific IPsec tunnel for the security association. | detail extensive |
| Local gateway | Gateway address of the local system. | All levels |
| Remote gateway | Gateway address of the remote system. | All levels |
| IPsec inside interface | Name of the logical interface hosting the IPsec tunnels. | All levels |
| Tunnel MTU | MTU of the IPsec tunnel. | All levels |

Table 79: show services ipsec-vpn ipsec security-associations Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------|--|-----------------|
| Local identity | <p>Protocol, address or prefix, and port number of the local entity of the IPsec association. The format is id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation). The protocol is always displayed as any because it is not user-configurable in the IPsec rule. Similarly, the port number field in the output is always displayed as 0 because it is not user-configurable in the IPsec rule. The value of the id-data-len parameter can be one of the following, depending on the address configured in the IPsec rule:</p> <ul style="list-style-type: none"> For an IPv4 address, the length is 4 and the value displayed is 3. For a subnet mask of an IPv4 address, the length is 8 and the value displayed is 7. For a range of IPv4 addresses, the length is 8 and the value displayed is 7. For an IPv6 address prefix, the length is 16 and the value displayed is 15. For a subnet mask of an IPv6 address prefix, the length is 32 and the value displayed is 31. For a range of IPv6 address prefixes, the length is 32 and the value displayed is 31. <p>The value of the id-data-presentation field denotes the IPv4 address or IPv6 prefix details. If the fully qualified domain name (FQDN) is specified instead of the address for the local peer of the IPsec association, it is displayed instead of the address details.</p> | All levels |
| Remote identity | <p>Protocol, address or prefix, and port number of the remote entity of the IPsec association. The format is id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation). The protocol is always displayed as any because it is not user-configurable in the IPsec rule. Similarly, the port number field in the output is always displayed as 0 because it is not user-configurable in the IPsec rule. The value of the id-data-len parameter can be one of the following, depending on the address configured in the IPsec rule:</p> <ul style="list-style-type: none"> For an IPv4 address, the length is 4 and the value displayed is 3. For a subnet mask of an IPv4 address, the length is 8 and the value displayed is 7. For a range of IPv4 addresses, the length is 8 and the value displayed is 7. For an IPv6 address prefix, the length is 16 and the value displayed is 15. For a subnet mask of an IPv6 address prefix, the length is 32 and the value displayed is 31. For a range of IPv6 address prefixes, the length is 32 and the value displayed is 31. <p>The value of the id-data-presentation field denotes the IPv4 address or IPv6 prefix details. If the fully qualified domain name (FQDN) is specified instead of the address for the remote peer of the IPsec association, it is displayed instead of the address details.</p> | All levels |
| Primary remote gateway | IP address of the configured primary remote peer. | All levels |
| Backup remote gateway | IP address of the configured backup remote peer. | All levels |

Table 79: show services ipsec-vpn ipsec security-associations Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|---|-------------------------|
| State | State of the primary or backup interface: Active , Offline , or Standby . Both ES PICs are initialized to Offline . For primary and backup peers, State can be Active or Standby . If both peers are in a state of Standby , no connection exists yet between the two peers. | All levels |
| Failover counter | Number of times a PIC switched between primary and backup interfaces, or the number of times the tunnel switched between the primary and remote peers since the software has been activated. | All levels |
| Direction | Direction of the security association: inbound or outbound . | All levels |
| SPI | Value of the security parameter index. | All levels |
| AUX-SPI | Value of the auxiliary security parameter index. <ul style="list-style-type: none"> When the value of Protocol is AH or ESP, AUX-SPI is always 0. When the value of Protocol is AH+ESP, AUX-SPI is always a positive integer. | All levels |
| Mode | Mode of the security association: <ul style="list-style-type: none"> transport—Protects single host-to-host protections. tunnel—Protects connections between security gateways. | detail extensive |
| Type | Type of security association: <ul style="list-style-type: none"> manual—Security parameters require no negotiation. They are static, and are configured by the user. dynamic—Security parameters are negotiated by the IKE protocol. Dynamic security associations are not supported in transport mode. | detail extensive |
| State | Status of the security association: <ul style="list-style-type: none"> Installed—The security association is installed in the security association database. (For transport mode security associations, the value of State must always be Installed.) Not installed—The security association is not installed in the security association database. | detail extensive |
| Protocol | Protocol supported: <ul style="list-style-type: none"> transport mode supports Encapsulation Security Protocol (ESP) or Authentication Header (AH). tunnel mode supports ESP or AH+ESP. | All levels |
| Authentication | Type of authentication used: hmac-md5-96 , hmac-sha1-96 , or none . | detail extensive |
| Encryption | Type of encryption algorithm used: aes-cbc (128 bits) , aes-cbc (192 bits) , aes-cbc (256 bits) , des-cbc , 3des-cbc , or None . | detail |

Table 79: show services ipsec-vpn ipsec security-associations Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|--|------------------|
| Soft lifetime
Hard lifetime | <p>Each lifetime of a security association has two display options, hard and soft, one of which must be present for a dynamic security association. The hard lifetime specifies the lifetime of the SA. The soft lifetime, which is derived from the hard lifetime, informs the IPsec key management system that the SA is about to expire. This information allows the key management system to negotiate a new SA before the hard lifetime expires.</p> <ul style="list-style-type: none"> • Expires in seconds seconds—Number of seconds left until the security association expires. • Expires in kilobytes kilobytes—Number of kilobytes left until the security association expires. | detail extensive |
| Anti-replay service | State of the service that prevents packets from being replayed: Enabled or Disabled . | detail extensive |
| Replay window size | Configured size, in packets, of the antireplay service window: 32 or 64 . The antireplay window size protects the receiver against replay attacks by rejecting old or duplicate packets. If the replay window size is 0 , antireplay service is disabled. | detail |

Sample Output

show services ipsec-vpn ipsec security associations extensive

```

user@host> show services ipsec-vpn ipsec security-associations extensive
Service set: service-set-1
  Rule: _junos_, Term: term-1, Tunnel index: 1
  Local gateway: 101.101.101.2, Remote gateway: 14.14.14.4
  IPSec inside interface: sp-2/0/0.1 Local identity:
  ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Primary remote gateway: 101.101.101.1, State: Standby
  Backup remote gateway: 14.14.14.4, State: Active
  Failover counter: 1

  Direction: inbound, SPI: 3743521590, AUX-SPI: 0
  Mode: tunnel, Type: dynamic, State: Installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
  Soft lifetime: Expires in 23043 seconds
  Hard lifetime: Expires in 23178 seconds
  Anti-replay service: Enabled, Replay window size: 64

  Direction: outbound, SPI: 2551045240, AUX-SPI: 0
  Mode: tunnel, Type: dynamic, State: Installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
  Soft lifetime: Expires in 23043 seconds
  Hard lifetime: Expires in 23178 seconds
  Anti-replay service: Enabled, Replay window size: 64

```

show services ipsec-vpn ipsec statistics

| | |
|---------------------------------|---|
| Syntax | show services ipsec-vpn ipsec statistics
<brief detail>
<remote-gw remote-peer-address>
<service-set service-set-name> |
| Release Information | Command introduced before Junos OS Release 7.4.
New fields added in Junos OS Release 10.0. |
| Description | (Adaptive services interface only) Display IPsec statistics for the specified service set. If no service set is specified, the statistics for all service sets are displayed. |
| Options | <p>none—Display standard IPsec statistics for all service sets.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>remote-gw remote-peer-address—(Optional) Display IPsec statistics for an individual IPsec tunnel and an individual remote host.</p> <p>service-set service-set-name—(Optional) Display information about a particular service set.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services ipsec-vpn ipsec statistics detail on page 1971
show services ipsec-vpn ipsec statistics remote-gw on page 1971 |
| Output Fields | Table 80 on page 1969 lists the output fields for the show services ipsec-vpn ipsec statistics command. Output fields are listed in the approximate order in which they appear. |

Table 80: show services ipsec-vpn ipsec statistics Output Fields

| Field Name | Field Description | Level of Output |
|-----------------------|---|-----------------|
| PIC | The physical interface on which the IPsec tunnel is configured. | All levels |
| Service set | Name of the service set for which the IPsec tunnel is defined. | All levels |
| Local gateway | Gateway address of the local system. | All levels |
| Remote gateway | Gateway address of the remote system. | All levels |
| Tunnel index | Numeric identifier of the specific IPsec tunnel for the security association. | All levels |

Table 80: show services ipsec-vpn ipsec statistics Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------|---|-----------------|
| ESP statistics | Encapsulation Security Payload (ESP) statistics: <ul style="list-style-type: none"> • Encrypted bytes—Total number of bytes encrypted by the local system across the IPsec tunnel. • Decrypted bytes—Total number of bytes decrypted by the local system across the IPsec tunnel. • Encrypted packets—Total number of packets encrypted by the local system across the IPsec tunnel. • Decrypted packets—Total number of packets decrypted by the local system across the IPsec tunnel. | All levels |
| AH Statistics | Authentication Header statistics: <ul style="list-style-type: none"> • Input bytes—Total number of bytes received by the local system across the IPsec tunnel. • Output bytes—Total number of bytes transmitted by the local system across the IPsec tunnel. • Input packets—Total number of packets received by the local system across the IPsec tunnel. • Output packets—Total number of packets transmitted by the local system across the IPsec tunnel. | All levels |
| Errors | <ul style="list-style-type: none"> • AH authentication failures—Number of authentication header (AH) failures. An AH failure occurs when there is a mismatch of the authentication header in a packet transmitted across an IPsec tunnel. • ESP authentication failures—Number of Encapsulation Security Payload (ESP) failures. An ESP failure occurs when there is an authentication mismatch in ESP packets. • ESP Decryption failures—Number of ESP decryption failures. • Bad headers—Number of invalid headers detected. • Bad trailers—Number of invalid trailers detected. • Replay before window drops—Number of replay errors. A replay error is generated when a duplicate packet is received within the replay window. • Replayed pkts—Number of packets replayed. • IP integrity errors—Number of IP integrity errors. • Exceeds tunnel MTU—Number of times the tunnel maximum transmission unit (MTU) value was exceeded. • Rule lookup failures—Number of rule lookup failures. • No SA errors—Number of errors resulting from a missing security association (SA). • Flow errors—Number of flow errors. • Misc errors—Number of miscellaneous errors. | All levels |

Sample Output

show services ipsec-vpn ipsec statistics detail

```

user@host> show services ipsec-vpn ipsec statistics
PIC: sp-0/2/0, Service set: ss0

ESP Statistics:
  Encrypted bytes:          0
  Decrypted bytes:         0
  Encrypted packets:       0
  Decrypted packets:       0
AH Statistics:
  Input bytes:             168
  Output bytes:            168
  Input packets:           2
  Output packets:          2
Errors:
  AH authentication failures: 0
  ESP authentication failures: 0
  ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
  Replay before window drops: 0, Replayed pkts: 0
  IP integrity errors: 0, Exceeds tunnel MTU: 0
  Rule lookup failures: 0, No SA errors: 0
  Flow errors: 0, Misc errors: 0

```

show services ipsec-vpn ipsec statistics remote-gw

```

user@host> show services ipsec-vpn ipsec statistics remote-gw 22.22.2.1
PIC: sp-3/1/0, Service set: service-set-2
Local gateway: 22.22.1.1, Remote gateway: 22.22.2.1, Tunnel index: 2
ESP Statistics:
  Encrypted bytes:          0
  Decrypted bytes:         0
  Encrypted packets:       0
  Decrypted packets:       0
AH Statistics:
  Input bytes:             0
  Output bytes:            0
  Input packets:           0
  Output packets:          0
Errors:
  AH authentication failures: 0
  ESP authentication failures: 0
  ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
  Replay before window drops: 0, Replayed pkts: 0
  IP integrity errors: 0, Exceeds tunnel MTU: 0
  Rule lookup failures: 0, No SA errors: 0
  Flow errors: 0, Misc errors: 0

```

show services link-services cpu-usage

| | |
|---------------------------------|--|
| Syntax | show services link-services cpu-usage
<brief detail>
<interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 8.4. |
| Description | (M Series and T Series routers only) Display information about Link Services IQ (LSQ) CPU usage. |
| Options | <p>none—Display standard information about CPU usage for all LSQ interfaces.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display information about the specified LSQ interface.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services link-services cpu-usage brief (AS PIC) on page 1974
show services link-services cpu-usage brief (MultiServices PIC) on page 1974
show services link-services cpu-usage detail (AS PIC) on page 1974
show services link-services cpu-usage detail (MultiServices PIC) on page 1975 |
| Output Fields | Table 81 on page 1972 lists the output fields for the show services link-services cpu-usage command. Output fields are listed in the approximate order in which they appear. |

Table 81: show services link-services cpu-usage Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------|---|-----------------|
| Role | CPU functional category. | brief |
| 1 Second Average | Percentage of usage during 1-second duration. | All levels |
| 5 Second Average | Percentage of usage during 5-second duration. | All levels |
| QoS | Quality of service (QoS) CPU, which takes care of queuing and scheduling of incoming IP packets on a per-bundle basis. It schedules packets with higher QoS values first. | All levels |
| Sequencer | Assigns sequence numbers to outgoing MLPPP fragments and interleaves link fragmentation and interleaving (LFI) traffic. | All levels |
| Load Balancer | Distributes load across different fragmenter CPUs. | All levels |
| Fragmenter | Main LSQ CPU; fragments IP packets into MLPPP fragments and also reassembles MLPPP fragments into IP packets. | All levels |
| Total | Sum of all CPU functions. | brief |

Table 81: show services link-services cpu-usage Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------------------|---|-----------------|
| Idle | Counts idle cycles when the CPU does not have any work. | detail |
| Timer | Takes care of periodic events driven by a timer, such as timeouts. | detail |
| System | System housekeeping thread. | detail |
| Input (QoS) | Acquires and queues incoming IP frames from hardware interfaces. | detail |
| Output (QoS) | Sends scheduled frames to the next processing CPU. | detail |
| Output Frags (QoS) | Sends outstanding frames to the fragmenter CPU. | detail |
| Bypass (QoS) | Sends outstanding frames for LFI. | detail |
| Free frame (QoS) | Frees dropped frames. | detail |
| CPUnumber | Identifier number of specific CPU. | detail |
| Drop (Fragmenter) | Drops frames that have been marked by the QoS CPU. | detail |
| Frag (Fragmenter) | Fragments IP frames into MLPPP fragments. | detail |
| Reass (Fragmenter) | Reassembles MLPPP fragments into IP frames. | detail |
| Freeback (Fragmenter) | Handles freeback of credits from other CPUs (MultiServices PICs only). | detail |
| Input LFI (Sequencer) | Receives LFI traffic from QoS CPU and transmits it with strict priority over MLPPP. | detail |
| Input Frag (Sequencer) | Receives MLPPP fragments from fragmenter CPUs, assigns sequence numbers, and appends MLPPP headers. | detail |
| Output Frag (Sequencer) | Load-balances and transmits fragments across links. | detail |
| Retry (Sequencer) | Retries transmission if hardware was busy in the previous attempt. | detail |
| Input Alloc (Load Balancer) | Acquires frames from hardware interfaces and validates them. | detail |
| Input (Load Balancer) | Performs error and sanity checks and check frames for PortMapping. | detail |
| Output (Load Balancer) | Sends frame to next processing CPU. | detail |

Table 81: show services link-services cpu-usage Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|--|-----------------|
| Freeback (Load Balancer) | Handles freeback of credits from other CPUs. | detail |

Sample Output

show services link-services cpu-usage brief (AS PIC)

```

user@host> show services link-services cpu-usage interface lsq-0/0/0 brief
Role           1 Second Average      5 Second Average
QoS              1.0%                    1.0%
Sequencer        0.1%                    0.1%
Fragmenter       0.1%                    0.1%
Total            0.1%                    0.1%

```

show services link-services cpu-usage brief (MultiServices PIC)

```

user@host> show services link-services cpu-usage interface lsq-0/0/0 brief
Role           1 Second Average      5 Second Average
QoS              0.1%                    0.1%
Fragmenter       0.1%                    0.1%
Load Balancer    0.0%                    0.0%
Total            0.1%                    0.1%

```

show services link-services cpu-usage detail (AS PIC)

```

user@host> show services link-services cpu-usage interface lsq-0/0/0 detail

QoS           Idle  Timer  System  Input  Output  Output  Bypass  Free
              frame
CPU0           99.1%  0.9%  0.0%  0.0%  0.0%  0.0%  0.0%  0.0%
CPU1           99.8%  0.1%  0.0%  0.0%  0.0%  0.0%  0.0%  0.0%
1 sec ave     99.5%  0.5%  0.0%  0.0%  0.0%  0.0%  0.0%  0.0%
5 sec ave     99.5%  0.5%  0.0%  0.0%  0.0%  0.0%  0.0%  0.0%

Fragmenter     Idle  Timer  System  Drop  Frag  Reass  Free
              back
CPU0           96.6%  0.1%  0.0%  0.0%  0.0%  3.3%  0.0%
CPU1           99.9%  0.1%  0.0%  0.0%  0.0%  0.0%  0.0%
CPU2           99.9%  0.1%  0.0%  0.0%  0.0%  0.0%  0.0%
CPU3           99.9%  0.1%  0.0%  0.0%  0.0%  0.0%  0.0%
CPU4           99.9%  0.1%  0.0%  0.0%  0.0%  0.0%  0.0%
CPU5           99.9%  0.1%  0.0%  0.0%  0.0%  0.0%  0.0%
CPU6           99.9%  0.1%  0.0%  0.0%  0.0%  0.0%  0.0%
CPU7           99.9%  0.1%  0.0%  0.0%  0.0%  0.0%  0.0%
CPU8           99.9%  0.1%  0.0%  0.0%  0.0%  0.0%  0.0%
1 sec ave     99.5%  0.1%  0.0%  0.0%  0.0%  0.4%  0.0%
5 sec ave     99.5%  0.1%  0.0%  0.0%  0.0%  0.4%  0.0%

Sequencer      Idle  System  Input  Input  Output  Retry
              LFI   Frag   Frag
CPU0           99.9%  0.1%  0.0%  0.0%  0.0%  0.0%
CPU1          100.0%  0.0%  0.0%  0.0%  0.0%  0.0%

```



```

1 sec ave      99.9%    0.1%   0.0%   0.0%   0.0%   0.0%
5 sec ave      99.9%    0.1%   0.0%   0.0%   0.0%   0.0%

```

show services link-services cpu-usage detail (MultiServices PIC)

```
user@host> show services link-services cpu-usage interface lsq-0/0/0 detail
```

| QoS | Idle | Timer | System | Input | Output | Output
Frag | Bypass | Free
frame |
|-----------|-------|-------|--------|-------|--------|----------------|--------|---------------|
| CPU0 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU1 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU2 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU3 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU4 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| 1 sec ave | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| 5 sec ave | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |

| Fragmenter | Idle | Timer | System | Drop | Frag | Reass | Free
back |
|------------|-------|-------|--------|------|------|-------|--------------|
| CPU0 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU1 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU2 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU3 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU4 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU5 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU6 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU7 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU8 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU9 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU10 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU11 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU12 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU13 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU14 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU15 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU16 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU17 | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| 1 sec ave | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| 5 sec ave | 99.9% | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |

| Load-Balancer | Idle | System | Input
Alloc | Input | Output | Free
back |
|---------------|--------|--------|----------------|-------|--------|--------------|
| CPU0 | 100.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| CPU1 | 100.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| 1 sec ave | 100.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |
| 5 sec ave | 100.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% |

show services l2tp multilink

| | |
|---------------------------------|---|
| Syntax | show services l2tp multilink
<brief detail extensive statistics>
<bundle-id <i>number</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M10i and M7i routers only) Display L2TP output organized by multilink bundle. |
| Options | <p>none—Same as brief.</p> <p>brief detail extensive statistics—(Optional) Display the specified level of output. Use the statistics option to display packets and bytes that have been encapsulated in the Multilink Protocol. Nonmultilink packets received on member sessions are not counted here.</p> <p>bundle-id <i>number</i>—(Optional) Display L2TP multilink bundle information for only the specified bundle.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • L2TP Services Configuration Overview on page 638 • L2TP Minimum Configuration on page 639 • clear services l2tp multilink on page 1829 |
| List of Sample Output | show services l2tp multilink extensive on page 1979 |
| Output Fields | Table 82 on page 1976 lists the output fields for the show services l2tp multilink command. Output fields are listed in the approximate order in which they appear. |

Table 82: show services l2tp multilink Output Fields

| Field Name | Field Description | Level of Output |
|------------------------|--|-----------------|
| Bundle ID | Bundle identifier. | All levels |
| Links | Number of links in the multilink bundle. | All levels |
| Bundle endpoint | Endpoint discriminator that represents the device transmitting the packet. | All levels |
| Input MRRU | Maximum packet size that the input interface can process. | detail |
| Output MRRU | Maximum packet size that the output interface can process. | detail |

Table 82: show services l2tp multilink Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------|--|-----------------|
| Session local ID | Identifier of the local endpoint of the L2TP session, as assigned by the L2TP network server (LNS). | detail |
| Session remote ID | Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC). | detail |
| State | Status of the L2TP session: <ul style="list-style-type: none"> • Established—The session is operating. • closed—The session is being closed. • destroyed—The session is being destroyed. • clean-up—The session is being cleaned up. • lns-ic-accept-new—A new session is being accepted. • lns-ic-idle—The session has been created and is idle. • lns-ic-reject-new—The new session is being rejected. • lns-ic-wait-connect—The session is waiting for the peer's incoming call connected (ICCN) message. | detail |
| Username | Name of the user logged in to the session. | detail |
| Mode | Mode of the interface representing the multilink bundle: dedicated or shared . | extensive |
| Local IP | IP address of the local endpoint of the Point-to-Point Protocol (PPP) session. | extensive |
| Remote IP | IP address of the remote endpoint of the PPP session. | extensive |
| Local name | Name of the LNS instance in which the session was created. | extensive |
| Remote name | Name of the LAC from which the session was created. | extensive |
| Local MRU | Maximum receive unit (MRU) setting of the local device, in bytes. | extensive |

Table 82: show services l2tp multilink Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|--|-----------------|
| Remote MRU | MRU setting of the remote device, in bytes. | extensive |
| Statistics since | <p>Date and time when collection of the following statistics began:</p> <ul style="list-style-type: none"> • Control Tx—Amount of control information transmitted, in packets and bytes. • Control Rx—Amount of control information received, in packets and bytes. • Data Tx—Amount of data transmitted, in packets and bytes. • Data Rx—Amount of data received, in packets and bytes. • Errors Tx—Number of errors transmitted, in packets. • Errors Rx—Number of errors received, in packets. • Lcp Echo Req Tx—Number of LCP echo requests transmitted, in packets. • Lcp Echo Req Rx—Number of LCP echo requests received, in packets. • Lcp Echo Rep Tx—Number of LCP echo responses transmitted, in packets. • Lcp Echo Rep Rx—Number of LCP echo responses received, in packets. • Lcp Echo Req Timeout—Number of LCP echo requests that timed out. • Lcp Echo Req Error—Number of errors received for LCP echo packets. • Lcp Echo Rep Error—Number of errors transmitted for LCP echo packets. • MRRU—Maximum packet size processed. • TX—Number of packets transmitted. • RX—Number of packets received. • link—Link of the multilink bundle associated with the L2TP session. | extensive |

Sample Output

show services l2tp multilink extensive

```

user@host> show services l2tp multilink extensive
Bundle ID: 1
  Links: 2, Bundle endpoint: user@juniper.com
  Input MRRU: 1524, Output MRRU: 1524
  Session local ID: 46122, Session remote ID: 39307
    State: Established, Username: user1@juniper.com, Mode: dedicated
    Local IP: 10.58.255.129:1701, Remote IP: 10.58.255.131:1701
    Local name: router3, Remote name: router4
  Session local ID: 4254, Session remote ID: 39308
    State: Established, Username: user2@juniper.com, Mode: dedicated
    Local IP: 10.1.255.1:1701, Remote IP: 10.1.255.2:1701
    Local name: router1, Remote name: router2
  Statistics since: Mon May 17 11:47:35 2004
    Packets      Bytes
    Control Tx   7      196
    Control Rx   3      90
    Data Tx      0       0
    Data Rx      0       0
    Errors Tx    0
    Errors Rx    0
    Lcp Echo Req Tx 0
    Lcp Echo Req Rx 0
    Lcp Echo Rep Tx 0
    Lcp Echo Rep Rx 0
    Lcp Echo Req Timeout 0
    Lcp Echo Req Error 0
    Lcp Echo Rep Error 0
  MRRU 1486 droptime 0 maxfrag 0 minfrag 32 minmru 1482 maxqlen 3000
    TX: Packets 0   Frags 0   Txseq 0x0
    RX: Packets 24  Frags 24  Rxseq 0x18  mseq 23  maxdiff 1  reass 24
      fragments copied 0
    link 0 : seq 0x17 mru 1482 encapslen 8 qlen 0 context 0xea01eb0

```

show services l2tp radius

| | |
|---------------------------------|---|
| Syntax | <pre>show services l2tp radius <accounting (servers statistics)> <authentication (servers statistics)> <servers> <statistics></pre> |
| Release Information | Command introduced in Junos OS Release 9.0. |
| Description | (M7i, M10i, and M120 routers only) Display RADIUS servers and statistics information for the RADIUS servers configured on the router. |
| Options | <p>You must include one of the following keywords to provide a valid completion for the command:</p> <p>accounting (servers statistics)—(Optional) Display RADIUS servers or statistical accounting information only.</p> <p>authentication (servers statistics)—(Optional) Display RADIUS servers or statistical authentication information only.</p> <p>servers—(Optional) Display RADIUS authentication and accounting server information only.</p> <p>statistics—(Optional) Display RADIUS authentication and accounting statistics information only.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • L2TP Services Configuration Overview on page 638 • L2TP Minimum Configuration on page 639 |
| List of Sample Output | show services l2tp radius servers on page 1982
show services l2tp radius statistics on page 1982 |
| Output Fields | <p>Table 83 on page 1980 lists the output fields for the show services l2tp radius command. Output fields are listed in the approximate order in which they appear.</p> |

Table 83: show services l2tp radius Output Fields

| Field Name | Field Description |
|-------------|--|
| IP Address | IP address of the server. |
| State | (servers keyword only) Present state of the server. |
| UDP Port | Number of the UDP port used to send authentication or accounting messages to the server. |
| Retry Count | (servers keyword only) Number of times the RADIUS client resends a packet if no ACK is received. |

Table 83: show services l2tp radius Output Fields (*continued*)

| Field Name | Field Description |
|----------------------------|---|
| Timeout | (servers keyword only) Length of time the client waits for an ACK before retransmission. |
| Pending Requests | (servers keyword only) Number of client pending authentication or accounting requests. |
| Maximum Sessions | (servers keyword only) Maximum number of pending requests on each RADIUS client before the server moves to the next RADIUS client, which is 200 times the maximum number of clients that can be created on a server (which is 12). |
| Dead Time | (servers keyword only) Interval to wait before retrying a server after it fails to send a response to an authentication or accounting request. |
| Secret Type | (servers keyword only) Secret type configured on the RADIUS server. |
| Profile | (servers keyword only) Name of profile configured for the RADIUS server. |
| Access requests | (statistics keyword only) Number of access requests sent to the server. |
| Rollover requests | (statistics keyword only) Number of requests coming into the server as a result of the previous server timing out. |
| Retransmissions | (statistics keyword only) Number of retransmissions. |
| Access accepts | (statistics keyword only) Number of access accept messages received from the server. |
| Access rejects | (statistics keyword only) Number of access reject messages received from the server. |
| Access challenges | (statistics keyword only) Number of access challenges received from the server. |
| Malformed responses | (statistics keyword only) Number of responses with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one). |
| Bad authenticators | (statistics keyword only) Number of responses in which the authenticator is incorrect for the matching request. This can occur if the RADIUS secrets for the client and server do not match. |
| Requests pending | (statistics keyword only) Number of requests waiting for a response. |
| Request timeouts | (statistics keyword only) Number of requests that timed out. |
| Unknown responses | (statistics keyword only) Number of unknown responses. The RADIUS response type in the header is invalid or unsupported. |
| Packets dropped | (statistics keyword only) Number of packets dropped because they are too short or because the router receives a response for which there is no corresponding request. For example, if the router sends a request that times out, the router removes the request from the list and sends a new request. If the server is slow and sends a response to the first request after the router removes the request, the packet is dropped. |

Sample Output

show services l2tp radius servers

```
user@host> show services l2tp radius servers
```

RADIUS Authentication Servers

| IP Address | State | UDP Port | Retry Count | Timeout | Pending Requests | Maximum Sessions | Dead Time | Secret Type |
|-----------------|--------|----------|-------------|---------|------------------|------------------|-----------|-------------|
| 17.1.1.1 | Active | 1812 | 2 | 25 | 0 | 2400 | 300 | radius-key |
| 133.122.1.1 | Active | 1812 | 5 | 35 | 0 | 2400 | 300 | radius-key |
| 134.141.1.1 | Active | 1812 | 2 | 25 | 0 | 2400 | 300 | radius-key |
| 172.28.30.174 | Active | 1812 | 7 | 75 | 0 | 2400 | 300 | radius-key |
| 172.28.30.175 | Active | 1812 | 7 | 75 | 0 | 2400 | 300 | radius-key |
| 172.28.30.176 | Active | 1812 | 4 | 55 | 0 | 2400 | 300 | radius-key |
| 172.128.30.176 | Active | 1812 | 3 | 3 | 0 | 2400 | 300 | none-set |
| 172.128.130.174 | Active | 1812 | 7 | 75 | 0 | 2400 | 300 | radius-key |

RADIUS Accounting Servers

| IP Address | State | UDP Port | Retry Count | Timeout | Pending Requests | Maximum Sessions | Dead Time | Secret Type |
|-----------------|--------|----------|-------------|---------|------------------|------------------|-----------|-------------|
| 17.1.1.1 | Active | 1813 | 2 | 25 | 0 | 2400 | 300 | radius-key |
| 133.122.1.1 | Active | 1813 | 5 | 35 | 0 | 2400 | 300 | radius-key |
| 134.141.1.1 | Active | 1813 | 2 | 25 | 0 | 2400 | 300 | radius-key |
| 172.28.30.174 | Active | 1813 | 7 | 75 | 0 | 2400 | 300 | radius-key |
| 172.28.30.175 | Active | 1813 | 7 | 75 | 0 | 2400 | 300 | radius-key |
| 172.28.30.176 | Active | 1813 | 4 | 55 | 0 | 2400 | 300 | radius-key |
| 172.128.30.176 | Active | 1813 | 3 | 3 | 0 | 2400 | 300 | none-set |
| 172.128.130.174 | Active | 1813 | 7 | 75 | 0 | 2400 | 300 | radius-key |

RADIUS Accounting Servers

```
Profile: user1
```

show services l2tp radius statistics

```
user@host> show services l2tp radius statistics
```

RADIUS Authentication Statistics

```
Authentication statistics:
```

```
Server 17.1.1.1, UDP port: 1812
```

```
Access requests      : 40
Rollover requests    : 5
Retransmissions      : 2
Access accepts       : 39
Access rejects       : 1
```



```
Access challenges      : 3
Malformed responses    : 0
Bad authenticators     : 0
Requests pending       : 1
Request timeouts       : 0
Unknown responses      : 0
Packets dropped        : 0
```

RADIUS Accounting Statistics

Accounting statistics:

Server 172.128.130.174, UDP port: 1813

```
Total requests        : 9
Start requests         : 6
Interim requests       : 1
Stop requests          : 2
Rollover requests      : 0
Retransmissions        : 1
Total response         : 9
Start responses        : 6
Interim responses      : 1
Stop responses         : 2
Malformed responses    : 0
Bad authenticators     : 0
Requests pending       : 1
Request timeouts       : 0
Unknown responses      : 0
Packets dropped        : 0
```

show services l2tp session

Syntax show services l2tp session
<brief | detail | extensive>
<interface *interface-name*>
<local-gateway *gateway-address*>
<local-gateway-name *gateway-name*>
<local-session-id *session-id*>
<local-tunnel-id *tunnel-id*>
<peer-gateway *gateway-address*>
<peer-gateway-name *gateway-name*>
<statistics>
<tunnel-group *group-name*>
<user *username*>

Release Information Command introduced before Junos OS Release 7.4.
Support for LAC on MX Series routers introduced in Junos OS Release 10.4.
Support for LNS on MX Series routers introduced in Junos OS Release 11.4.

Description (M10i and M7i routers only) Display information about active L2TP sessions for LNS.

(MX Series routers only) Display information about active L2TP sessions for LAC and LNS.

Options **none**—Display standard information about all active L2TP sessions.

brief | detail | extensive—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display L2TP session information for only the specified adaptive services or inline services interface. The interface type depends on the line card as follows:

- **si-*fpc/pic/port***—MPCs on MX Series routers only. This option is not available for L2TP on M Series routers.
- **sp-*fpc/pic/port***—AS or Multiservices PICs on M7i, M10i, and M120 routers only. This option is not available for L2TP on MX Series routers.

local-gateway *gateway-address*—(Optional) Display L2TP session information for only the specified local gateway address.

local-gateway-name *gateway-name*—(Optional) Display L2TP session information for only the specified local gateway name.

local-session-id *session-id*—(Optional) Display L2TP session information for only the specified local session identifier.

local-tunnel-id *tunnel-id*—(Optional) Display L2TP session information for only the specified local tunnel identifier.

peer-gateway *gateway-address*—(Optional) Display L2TP session information for only the specified peer gateway address.

peer-gateway-name *gateway-name*—(Optional) Display L2TP session information for only the specified peer gateway name.

statistics—(Optional) Display the number of control packets and bytes transmitted and received for the session. You cannot include this option with any of the level options, **brief**, **detail**, or **extensive**.

tunnel-group *group-name*—(Optional) Display L2TP session information for only the specified tunnel group. To display information about L2TP CPU and memory usage, you can include the tunnel group name in the **show services service-sets memory-usage *group-name*** and **show services service-sets cpu-usage *group-name*** commands. This option is not available for L2TP LAC on MX Series routers.

user *username*—(M Series routers only) (Optional) Display L2TP session information for only the specified username.

Required Privilege Level view

Related Documentation

- [L2TP Services Configuration Overview on page 638](#)
- [L2TP Minimum Configuration on page 639](#)
- [clear services l2tp session on page 1830](#)

List of Sample Output

[show services l2tp session \(LNS on M Series Routers\) on page 1988](#)
[show services l2tp session \(LNS on MX Series Routers\) on page 1989](#)
[show services l2tp session \(LAC\) on page 1989](#)
[show services l2tp session detail \(LAC\) on page 1989](#)
[show services l2tp session extensive \(LAC\) on page 1989](#)
[show services l2tp session extensive \(LAC on MX Series Routers\) on page 1989](#)
[show services l2tp session extensive \(LNS on M Series Routers\) on page 1990](#)
[show services l2tp session extensive \(LNS on MX Series Routers\) on page 1990](#)
[show services l2tp session statistics \(MX Series Routers\) on page 1991](#)

Output Fields [Table 84 on page 1985](#) lists the output fields for the **show services l2tp session** command. Output fields are listed in the approximate order in which they appear.

Table 84: show services l2tp session Output Fields

| Field Name | Field Description | Level of Output |
|------------------|---|-----------------|
| Interface | (LNS only) Name of an adaptive services interface. | All levels |
| Tunnel group | (LNS only) Name of a tunnel group. | All levels |
| Tunnel local ID | Identifier of the local endpoint of the tunnel, as assigned by the L2TP network server (LNS). | All levels |
| Session local ID | Identifier of the local endpoint of the L2TP session, as assigned by the LNS. | All levels |

Table 84: show services l2tp session Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------------|---|------------------|
| Session remote ID | Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC). | All levels |
| State | State of the L2TP session: <ul style="list-style-type: none"> • Established—Session is operating. This is the only state supported for the LAC. • closed—Session is being closed. • destroyed—Session is being destroyed. • clean-up—Session is being cleaned up. • lns-ic-accept-new—New session is being accepted. • lns-ic-idle—Session has been created and is idle. • lns-ic-reject-new—New session is being rejected. • lns-ic-wait-connect—Session is waiting for the peer's incoming call connected (ICCN) message. | All levels |
| Bundle ID | (LNS only) Bundle identifier. Indicates the session is part of a multilink bundle. Sessions that have a blank Bundle field are not participating in the Multilink Protocol. Sessions in a multilink bundle might belong to different L2TP tunnels. For L2TP output organized by bundle ID, issue the show services l2tp multilink extensive command. | All levels |
| Mode | (LNS) Mode of the interface representing the session: shared or exclusive .

(LAC) Mode of the interface representing the session: shared or dedicated . Only dedicated is currently supported for the LAC. | extensive |
| Local IP | IP address of local endpoint of the Point-to-Point Protocol (PPP) session. | extensive |
| Remote IP | IP address of remote endpoint of the PPP session. | extensive |
| Username | (LNS only) Name of the user logged in to the session. | All levels |
| Assigned IP address | (LNS only) IP address assigned to remote client. | extensive |
| Local name | For LNS, name of the LNS instance in which the session was created. For LAC, name of the LAC. | extensive |
| Remote name | For LNS, name of the LAC from which the session was created. For LAC, name of the LAC instance. | extensive |
| Local MRU | (LNS only) Maximum receive unit (MRU) setting of the local device, in bytes. | extensive |
| Remote MRU | (LNS only) MRU setting of the remote device, in bytes. | extensive |

Table 84: show services l2tp session Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|---|------------------|
| Tx speed | <p>Transmit speed of the session conveyed from the LAC to the LNS, in bits per second (bps).</p> <p>Either the initial (initial) line speed or both the initial and current (update) line speeds can be displayed on MX Series routers:</p> <ul style="list-style-type: none"> • When connection speed updates are not enabled, then only the initial line speed is displayed. • When connection speed updates are enabled, then both the initial and the current speeds are displayed. <p>When the Tx connect speed method is set to none, the value of zero (0) is displayed.</p> | extensive |
| Rx speed | <p>Receive speed of the session conveyed from the LAC to the LNS, in bits per second (bps).</p> <p>Either the initial (initial) line speed or both the initial and current (update) line speeds can be displayed on MX Series routers:</p> <ul style="list-style-type: none"> • When connection speed updates are not enabled, then only the initial line speed is displayed. • When connection speed updates are enabled, then both the initial and the current speeds are displayed. <p>When the Tx connect speed method is set to none, the value of zero (0) is displayed.</p> | extensive |
| Bearer type | <p>Type of bearer enabled:</p> <ul style="list-style-type: none"> • 0—Might indicate that the call was not received over a physical link (for example, when the LAC and PPP are located in the same subsystem). • 1—Digital access requested. • 2—Analog access requested. • 4—Asynchronous Transfer Mode (ATM) bearer support. | extensive |
| Framing type | <p>Type of framing enabled:</p> <ul style="list-style-type: none"> • 1—Synchronous framing • 2—Asynchronous framing | extensive |
| LCP renegotiation | <p>(LNS only) Whether Link Control Protocol (LCP) renegotiation is configured: On or Off.</p> | extensive |
| Authentication | <p>Type of authentication algorithm used: Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).</p> | extensive |
| Interface ID | <p>(LNS only) Identifier used to look up the logical interface for this session.</p> | extensive |
| Interface unit | <p>Logical interface for this session.</p> | All levels |
| Call serial number | <p>Unique serial number assigned to the call.</p> | extensive |

Table 84: show services l2tp session Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------------|--|------------------|
| Policer bandwidth | Maximum policer bandwidth configured for this session. | extensive |
| Policer burst size | Maximum policer burst size configured for this session. | extensive |
| Firewall filter | Configured firewall filter name. | extensive |
| Session encapsulation overhead | Overhead allowance configured for this session, in bytes. | extensive |
| Session cell overhead | Cell overhead activation (On or Off). | extensive |
| Create time | Date and time when the call was created. | extensive |
| Up time | Length of time elapsed since the call became active, in hours, minutes, and seconds. | extensive |
| Idle time | Length of time elapsed since the call became idle, in hours, minutes, and seconds. | extensive |
| Statistics since | Date and time when collection of the following statistics began: <ul style="list-style-type: none"> • Control Tx—Amount of control information transmitted, in packets and bytes. • Control Rx—Amount of control information received, in packets and bytes. • Data Tx—Amount of data transmitted, in packets and bytes. • Data Rx—Amount of data received, in packets and bytes. • Errors Tx—Number of errors transmitted, in packets. • Errors Rx—Number of errors received, in packets. • LCP echo req Tx—Number of LCP echo requests transmitted, in packets. • LCP echo req Rx—Number of LCP echo requests received, in packets. • LCP echo rep Tx—Number of LCP echo responses transmitted, in packets. • LCP echo rep Rx—Number of LCP echo responses received, in packets. • LCP echo Req timeout—Number of LCP echo requests that timed out. • LCP echo Req error—Number of errors received for LCP echo packets. • LCP echo Rep error—Number of errors transmitted for LCP echo packets. | extensive |

Sample Output

show services l2tp session (LNS on M Series Routers)

```

user@host> show services l2tp session
Interface: sp-1/2/0, Tunnel group: group1, Tunnel local ID: 8802
  Local Remote Interface State      Bundle Username
  ID   ID   unit
  37966      5      2 Established

```

show services l2tp session (LNS on MX Series Routers)

```

user@host> show services l2tp session
Tunnel local ID: 40553
  Local Remote State Interface Interface
  ID ID unit Name
  17967 1 Established 1073749824 si-5/2/0

```

show services l2tp session (LAC)

```

user@host> show services l2tp session
Tunnel local ID: 31889
  Local Remote State Interface Interface
  ID ID unit Name
  31694 1 Established 311 pp0

```

show services l2tp session detail (LAC)

```

user@host> show services l2tp session detail
Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID: 1, Interface unit: 311
  State: Established, Interface: pp0, Mode: Dedicated
  Local IP: 10.1.1.2:1701, Remote IP: 10.1.1.1:1701
  Local name: ce-lac, Remote name: ce-lns

```

show services l2tp session extensive (LAC)

```

user@host> show services l2tp session extensive
Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID: 1
  Interface unit: 311
  State: Established, Mode: Dedicated
  Local IP: 10.10.1.2:1701, Remote IP: 10.10.1.1:1701
  Local name: ce-lac, Remote name: ce-lns
  Tx speed: 0, Rx speed: 0
  Bearer type: 1, Framing type: 1
  LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
  Interface unit: 311, Call serial number: 0
  Policer bandwidth: 0, Policer burst size: 0
  Policer exclude bandwidth: 0, Firewall filter: 0
  Session encapsulation overhead: 0, Session cell overhead: 0
  Create time: Tue Aug 24 14:38:23 2010, Up time: 01:06:25
  Idle time: N/A

```

show services l2tp session extensive (LAC on MX Series Routers)

```

user@host> show services l2tp session extensive
Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID: 1
  Interface unit: 311
  State: Established, Mode: Dedicated
  Local IP: 10.10.1.2:1701, Remote IP: 10.10.1.1:1701
  Local name: ce-lac, Remote name: ce-lns
  Tx speed: initial 64000, Update 256000
  Rx speed: initial 64000, Update 256000
  Bearer type: 1, Framing type: 1
  LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
  Interface unit: 311, Call serial number: 0
  Policer bandwidth: 0, Policer burst size: 0
  Policer exclude bandwidth: 0, Firewall filter: 0
  Session encapsulation overhead: 0, Session cell overhead: 0

```

Create time: Tue Aug 24 14:38:23 2010, Up time: 01:06:25
Idle time: N/A

show services l2tp session extensive (LNS on M Series Routers)

```

user@host> show services l2tp session extensive
Interface: sp-1/2/0, Tunnel group: group1, Tunnel local ID: 62746
Session local ID: 56793, Session remote ID: 53304
State: Established, Bundle ID: 5, Mode: shared
Local IP: 10.128.1.1:1701, Remote IP: 10.128.1.2:1701
Username: usr1@juniper_1.net, Assigned IP address: 10.50.2.1/32
Local MRU: 4000, Remote MRU: 1500, Tx speed: 64000, Rx speed: 64000
Bearer type: 2, Framing type: 1
LCP renegotiation: Off, Authentication: CHAP, Interface ID: unit_20
Interface unit: 20, Call serial number: 4137941434
Policer bandwidth: 64000, Policer burst size: 51200
Firewall filter: f1
Session encapsulation overhead: 16, Session cell overhead: On
Create time: Tue Mar 23 14:13:15 2004, Up time: 01:16:41
Idle time: 00:00:00
Statistics since: Tue Mar 23 14:13:13 2004

```

| | Packets | Bytes |
|------------|---------|-------|
| Control Tx | 4 | 88 |
| Control Rx | 2 | 28 |
| Data Tx | 0 | 0 |
| Data Rx | 461 | 29.0k |
| Errors Tx | 0 | |
| Errors Rx | 0 | |

```

Interface: sp-1/2/0, Tunnel group: group_company_dns, Tunnel local ID: 37266
Session local ID: 39962, Session remote ID: 53303
State: Established, Bundle ID: 5, Mode: shared
Local IP: 10.128.11.1:1701, Remote IP: 10.128.11.2:1701
Username: usr1@company.com, Assigned IP address: 10.46.2.3/24
Local name: router-1, Remote name: router-2
Local MRU: 4470, Remote MRU: 4470, Tx speed: 155000000, Rx speed: 155000000
Bearer type: 2, Framing type: 1
LCP renegotiation: Off, Authentication: CHAP, Interface ID: unit_31
Interface unit: 31, Call serial number: 4137941433
Policer bandwidth: 64000, Policer burst size: 51200
Firewall filter: f1
Create time: Tue Mar 23 14:13:17 2004, Up time: 01:16:39
Idle time: 01:16:36
Statistics since: Tue Mar 23 14:13:15 2004

```

| | Packets | Bytes |
|------------|---------|-------|
| Control Tx | 6 | 196 |
| Control Rx | 4 | 150 |
| Data Tx | 0 | 0 |
| Data Rx | 1 | 80 |
| Errors Tx | 0 | |
| Errors Rx | 0 | |

show services l2tp session extensive (LNS on MX Series Routers)

```

user@host> show services l2tp session extensive
Tunnel local ID: 40553
Session local ID: 17967, Session remote ID: 1
Interface unit: 1073749824
State: Established
Interface: si-5/2/0
Mode: Dedicated

```



```

Local IP: 11.1.1.2:1701, Remote IP: 11.1.1.3:1701
Local name: lns-mx960, Remote name: testlac
Tx speed: 56000, Rx speed: 0
Bearer type: 2, Framing type: 1
LCP renegotiation: Off, Authentication: None
Call serial number: 1
Create time: Mon Apr 25 20:27:50 2011, Up time: 00:01:48
Idle time: N/A
Statistics since: Mon Apr 25 20:27:50 2011

```

| | Packets | Bytes |
|------------|---------|-------|
| Control Tx | 4 | 219 |
| Control Rx | 4 | 221 |
| Data Tx | 0 | 0 |
| Data Rx | 10 | 228 |
| Errors Tx | 0 | |
| Errors Rx | 0 | |

show services l2tp session statistics (MX Series Routers)

```

user@host>show services l2tp session statistics local session-id 1
Tunnel local ID: 17185
Session local ID: 1, Session remote ID: 14444, Interface unit: 1073788352
State: Established
Statistics since: Mon Aug 1 13:27:47 2011

```

| | Packets | Bytes |
|---------|---------|-------|
| Data Tx | 4 | 51 |
| Data Rx | 3 | 36 |

show services l2tp summary

| | |
|---------------------------------|---|
| Syntax | <code>show services l2tp summary</code>
<code><interface sp-fpc/pic/port></code>
<code><statistics></code> |
| Release Information | Command introduced before Junos OS Release 7.4.
Support for LAC on MX Series routers introduced in Junos OS Release 10.4.
Support for LNS on MX Series routers introduced in Junos OS Release 11.4.
Support for statistics option introduced in Junos OS Release 13.1. |
| Description | (M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Display Layer 2 Tunneling Protocol (L2TP) summary information. |
| Options | <p>none—Display complete L2TP summary information. For LNS on M Series routers, display L2TP summary information for all adaptive services interfaces. For LNS on MX Series routers, display L2TP summary information for all inline services interfaces.</p> <p>interface sp-fpc/pic/port—(Optional) Display L2TP summary information for only the specified adaptive services interface. This option is not available for L2TP on MX Series routers.</p> <p>statistics—(Optional) Display a summary of control packets and bytes transmitted and received.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> L2TP Services Configuration Overview on page 638 L2TP Minimum Configuration on page 639 |
| List of Sample Output | show services l2tp summary (LAC on M Series routers) on page 1995
show services l2tp summary (LAC on MX Series routers) on page 1995
show services l2tp summary (LNS on MX Series routers) on page 1996
show services l2tp summary (LNS on M Series routers) on page 1996
show services l2tp summary statistics (MX Series routers) on page 1996 |
| Output Fields | Table 85 on page 1992 lists the output fields for the show services l2tp summary command. Output fields are listed in the approximate order in which they appear. |

Table 85: show services l2tp summary Output Fields

| Field Name | Field Description |
|---|---|
| Failover within a preference level | State of this tunnel selection method on the LAC. When enabled, tunnel selection fails over within a preference level. When disabled, tunnel selection drops to the next lower preference level. Not displayed for LNS on M Series routers. |

Table 85: show services l2tp summary Output Fields (*continued*)

| Field Name | Field Description |
|--|---|
| Weighted load balancing | State of this tunnel selection method on the LAC. When enabled, the maximum session limit of a tunnel determines its weight within a preference level. Tunnel selection proceeds from greatest to least weight. When disabled, selection defaults to a round robin method. Not displayed for LNS on M Series routers. |
| Destination equal load balancing | State of this tunnel selection method on the LAC. When enabled, the LAC selects tunnels based on the session count for destinations and the tunnel session count. Not displayed for LNS on M Series routers. |
| Tunnel authentication challenge | State of tunnel authentication, indicating whether the LAC and LNS exchange an authentication challenge and response during the establishment of the tunnel. The state is Enabled when a secret is configured in the tunnel profile or on the RADIUS server in the Tunnel-Password attribute [69]. The state is Disabled when the secret is not present. Not displayed for LNS on M Series routers. |
| Calling number avp | When the state is Enabled , the LAC includes the value of the Calling Number AVP 22 in ICRQ packets sent to the LNS. When the state is Disabled , the attribute is not sent to the LNS. Not displayed for LNS on M Series routers. |
| Failover Protocol | When the state is enabled, the LAC operates in the default <i>failover-protocol-fall-back-to-silent-failover</i> manner. When the state is disabled, the disable-failover-protocol statement has been issued and the LAC operates only in silent failover mode. Not displayed for LNS on M Series routers. |
| Tx connect speed method | <p>The connection speed method configured to send the speed values in the L2TP Tx Connect Speed (AVP 24) and L2TP Rx Connect Speed (AVP 38). Possible values are:</p> <ul style="list-style-type: none"> • actual
This is the default value. • ancp • none • pppoe-ia-tag • static |
| Rx speed avp when equal | Indicates if the Rx connect speed when equal configuration is enabled or disabled . |

Table 85: show services l2tp summary Output Fields (*continued*)

| Field Name | Field Description |
|---|--|
| Tunnel assignment id | <p>Format of the tunnel name.</p> <p>Format of the tunnel name, based on RADIUS attributes returned from the AAA server:</p> <ul style="list-style-type: none"> • authentication-id—Name consists of only Tunnel Assignment-Id [82]. This is the default value. • client-server-id—Name is a combination of Tunnel-Client-Auth-Id [90], Tunnel-Server-Endpoint [67], and Tunnel-Assignment-Id [82]. This format is available only on MX Series routers. |
| Tunnel Tx Address Change | <p>Action taken by LAC when it receives a request from a peer to change the destination IP address, UDP port, or both:</p> <ul style="list-style-type: none"> • accept—Accepts change requests for the IP address or UDP port. This is the default action. • ignore—Ignores all change requests. • ignore-ip-address—Ignores change requests for the IP address but accepts them for the UDP port. • ignore-udp-port—Ignores change requests for the UDP port but accepts them for the IP address. |
| Min Retransmission Timeout for control packets | Minimum number of seconds that the local peer waits for the initial response after transmitting an L2TP control packet. If no response has been received by the time the period expires, the local peer retransmits the packet. |
| Min Retransmission Timeout for control packets | Minimum number of seconds that the local peer waits for the initial response after transmitting an L2TP control packet. If no response has been received by the time the period expires, the local peer retransmits the packet. |
| Max Retransmissions for Established Tunnel | Maximum number of times control messages are retransmitted for established tunnels. |
| Max Retransmissions for Not Established Tunnel | Maximum number of times control messages are retransmitted for tunnels that are not established. |
| Tunnel Idle Timeout | Period that a tunnel can be inactive—that is, carrying no traffic—before it times out and is torn down. |
| Destruct Timeout | Period that the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed. |
| Reassembly Service Set | Indicates active IP reassembly configured for the interface. |
| Destination Lockout Timeout | Timeout period for which all future destinations are locked out, meaning that they are not considered for selection when a new tunnel is created. |

Table 85: show services l2tp summary Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------|--|
| Access Line Information | State of LAC global configuration for forwarding subscriber line information to the LNS, Enabled or Disabled .

Indicates active IP reassembly configured for the interface. |
| Speed Updates | State of LAC global configuration for including connection speed updates when it forwards subscriber line information to the LNS, Enabled or Disabled . |
| Destinations | Number of L2TP destinations for the LAC. Not displayed for LNS on M Series routers. |
| Tunnels | Number of L2TP tunnels established on the router. |
| Sessions | Number of L2TP sessions established on the router. |
| Switched sessions | Number of L2TP tunnel-switched sessions established on the router. |
| Control | Count of L2TP control packets and bytes sent and received. |
| Data | Count of L2TP data packets and bytes sent and received. |
| Errors | Count of L2TP error packets and bytes sent and received. |

Sample Output

show services l2tp summary (LAC on M Series routers)

```

user@host> show services l2tp summary
Administrative state is Drain
Failover within a preference level is Disabled
Weighted load balancing is Enabled
Destination equal load balancing is Disabled
Tunnel authentication challenge is Enabled
Calling number avp is Enabled
Failover Protocol is Disabled
Tunnel assignment id format is authentication-id
Destinations: 1 Tunnels: 1, Sessions: 1
  Tx packets    Rx packets    Memory (bytes)
Control      260           144          11513856
Data         7.5k          16.9k          8.3k
Errors         0             0

```

show services l2tp summary (LAC on MX Series routers)

```

user@host> show services l2tp summary
Administrative state is Drain
Failover within a preference level is Disabled
Weighted load balancing is Disabled
Destination equal load balancing is Enabled
Tunnel authentication challenge is Enabled
Calling number avp is Enabled

```

```
Failover Protocol is Disabled
Tx Connect speed method is static
Rx speed avp when equal is enabled
Tunnel Tx Address Change is Accept
Min Retransmissions Timeout for control packets is 2 seconds
Max Retransmissions for Established Tunnel is 7
Max Retransmissions for Not Established Tunnel is 5
Tunnel Idle Timeout is 60 seconds
Destruct Timeout is 300 seconds
Destination Lockout Timeout is 300 seconds
Reassembly Service Set is ssnr3
Destinations: 0, Tunnels: 0, Sessions: 0, Switched sessions: 0
```

show services l2tp summary (LNS on MX Series routers)

```
user@host show services l2tp summary
Administrative state is Drain
Failover within a preference level is Disabled
Weighted load balancing is Disabled
Destination equal load balancing is Disabled
Tunnel authentication challenge is Enabled
Calling number avp is Enabled
Failover Protocol is Enabled
Tx Connect speed method is static
reassembly Service Set is ssnr3
Destinations: 4, Tunnels: 19, Sessions: 65, Switched sessions: 2
```

show services l2tp summary (LNS on M Series routers)

```
user@host> show services l2tp summary
Tunnels: 2, Sessions: 2, Errors: 0
  Tx packets  Rx packets  Memory (bytes)
Control      6k          9k          688k
Data        70k         70k         3054
```

show services l2tp summary statistics (MX Series routers)

```
user@host>show services l2tp summary statistics
Administrative state is Drain
Failover within a preference level is Disabled
Weighted load balancing is Disabled
Destination equal load balancing is Disabled
Tunnel authentication challenge is Enabled
Calling number avp is Enabled
Failover Protocol is Enabled
Tx Connect speed method is advisory
Tunnel assignment id format is assignment-id
Tunnel Tx Address Change is Accept
Min Retransmissions Timeout for control packets is 4 seconds
Max Retransmissions for Established Tunnel is 7
Max Retransmissions for Not Established Tunnel is 5
Tunnel Idle Timeout is 60 seconds
Destruct Timeout is 300 seconds
Destination Lockout Timeout is 300 secondsDestinations: 1, Tunnels: 1, Sessions:
31815, Switched sessions: 0
  Tx packets  Rx packets  Memory (bytes)
Control      90.4k      32.0k      245678080
Data        127.3k    100.8kk      0
Errors         0         0
```

show services l2tp tunnel

| | |
|---------------------------------|---|
| Syntax | <pre>show services l2tp tunnel <brief detail extensive> <interface sp-fpc/pic/port> <local-gateway gateway-address> <local-gateway-name gateway-name> <local-tunnel-id tunnel-id> <peer-gateway gateway-address> <peer-gateway-name gateway-name> <statistics> <tunnel-group group-name></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | <p>(M10i and M7i routers only) Display information about active Layer 2 Tunneling Protocol (L2TP) tunnels for LNS.</p> <p>(MX Series routers only) Display information about L2TP tunnels for LAC and LNS.</p> |
| Options | <p>none—Display standard information about all active L2TP tunnels.</p> <p>brief detail extensive—(Default) Display the specified level of output.</p> <p>interface sp-fpc/pic/port—(Optional) Display L2TP tunnel information for only the specified adaptive services interface. This option is not available for L2TP on MX Series routers.</p> <p>local-gateway gateway-address—(Optional) Display L2TP tunnel information for only the specified local gateway address.</p> <p>local-gateway-name gateway-name—(Optional) Display L2TP tunnel information for only the specified local gateway name.</p> <p>local-tunnel-id tunnel-id—(Optional) Display L2TP tunnel information for only the specified local tunnel identifier.</p> <p>peer-gateway gateway-address—(Optional) Display L2TP tunnel information for only the specified peer gateway address.</p> <p>peer-gateway-name gateway-name—(Optional) Display L2TP tunnel information for only the specified peer gateway name.</p> <p>statistics—(Optional) Display the number of control packets and bytes transmitted and received for the tunnel. You cannot include this option with any of the level options, brief, detail, or extensive.</p> <p>tunnel-group group-name—(Optional) Display L2TP tunnel information for only the specified tunnel group.</p> |
| Required Privilege Level | view |

- Related Documentation**
- [L2TP Services Configuration Overview on page 638](#)
 - [L2TP Minimum Configuration on page 639](#)

- List of Sample Output**
- [show services l2tp tunnel \(LAC\) on page 2000](#)
 - [show services l2tp tunnel detail \(LAC\) on page 2000](#)
 - [show services l2tp tunnel detail \(LAC on MX Series Routers\) on page 2000](#)
 - [show services l2tp tunnel detail \(LNS on MX Series Routers\) on page 2000](#)
 - [show services l2tp tunnel extensive \(LAC\) on page 2001](#)
 - [show services l2tp tunnel extensive \(LNS on M Series Routers\) on page 2001](#)
 - [show services l2tp tunnel extensive \(LNS on MX Series Routers\) on page 2002](#)
 - [show services l2tp tunnel statistics \(MX Series Routers\) on page 2002](#)

- Output Fields** [Table 86 on page 1998](#) lists the output fields for the **show services l2tp tunnel** command. Output fields are listed in the approximate order in which they appear.

Table 86: show services l2tp tunnel Output Fields

| Field Name | Field Description |
|---------------------|---|
| Interface | (LNS only) Name of an adaptive services interface. |
| Tunnel group | (LNS only) Name of a tunnel group. |
| Local ID | <p>On the LNS, number assigned by the LNS that identifies the local endpoint of the tunnel relative to the LNS: the LNS.</p> <p>On the LAC, number assigned by the LAC that identifies the local endpoint of the tunnel relative to the LAC: the LAC.</p> |
| Remote ID | <p>On the LNS, number assigned by the LAC that identifies the remote endpoint of the tunnel relative to the LNS: the LAC.</p> <p>On the LAC, number assigned by the LNS that identifies the remote endpoint of the tunnel relative to the LAC: the LNS.</p> |
| Remote IP | IP address of the peer endpoint of the tunnel. |
| Sessions | Number of L2TP sessions established through the tunnel. |

Table 86: show services l2tp tunnel Output Fields (*continued*)

| Field Name | Field Description |
|--|--|
| State | <p>State of the L2TP tunnel:</p> <ul style="list-style-type: none"> • cc_responder_accept_new—The tunnel has received and accepted the start control connection request (SCCRQ). • cc_responder_reject_new—The tunnel has received and rejected the SCCRQ. • cc_responder_idle—The tunnel has just been created. • cc_responder_wait_ctl_conn—The tunnel has sent the start control connection response (SCCRP) and is waiting for the start control connection connected (SCCCN) message. • clean-up—The tunnel is being cleaned up. • closed—The tunnel is being closed. • destroyed—The tunnel is being destroyed. • Established—The tunnel is operating. This is the only state supported for the LAC. • Terminate—The tunnel is terminating. • Unknown—The tunnel is not connected to the router. |
| Tunnel Name | (LAC only) Name of the created tunnel. This value includes the destination name followed by the value of the RADIUS Tunnel-Assignment-ID VSA [82]. |
| Local IP | IP address of the local endpoint of the tunnel. |
| Local name | Name used for local tunnel endpoint during tunnel negotiation. |
| Remote name | Name used for remote tunnel endpoint during tunnel negotiation. |
| Effective Peer Resync Mechanism | <p>(LAC only) Peer resynchronization mechanism (PRM) in effect for the tunnel:</p> <ul style="list-style-type: none"> • Failover protocol • Silent failover—Recovery takes place in the failed endpoint only using the proprietary silent failover protocol. |
| Nas Port Method | <p>NAS port method (type), which indicates whether the LAC sends Cisco NAS Port Info AVP (100) in ICRQs to the LNS:</p> <ul style="list-style-type: none"> • cisco-avp—sends the AVP. • none—does not send the AVP. |
| Tunnel Logical System | Logical system in which the L2TP tunnel is brought up. |
| Tunnel Routing Instance | Routing instance in which the L2TP tunnel is brought up. |
| Max sessions | Maximum number of sessions that can be established on this tunnel. |
| Window size | Number of control messages that can be sent without receipt of an acknowledgment. |
| Hello interval | Interval between the transmission of hello messages, in seconds. |

Table 86: show services l2tp tunnel Output Fields (*continued*)

| Field Name | Field Description |
|-------------------------|---|
| Create time | Date and time when the tunnel was created. While the LNS and LAC are connected, this value should correspond to the router's uptime. If connection to the LAC is severed, the State changes to Unknown and the Create time value resets. |
| Up time | Amount of time elapsed since the tunnel became active, in hours, minutes, and seconds. |
| Idle time | Amount of time elapsed since the tunnel became idle, in hours, minutes, and seconds. |
| Statistics since | <p>Date and time when collection of the following statistics began:</p> <ul style="list-style-type: none"> • Control Tx—Amount of control information transmitted, in packets and bytes. • Control Rx—Amount of control information received, in packets and bytes. • Data Tx—Amount of data transmitted, in packets and bytes. • Data Rx—Amount of data received, in packets and bytes. • Errors Tx—Number of errors transmitted, in packets. • Errors Rx—Number of errors received, in packets. |

Sample Output

show services l2tp tunnel (LAC)

```

user@host> show services l2tp tunnel
      Local ID  Remote ID  Remote IP          Sessions  State
          17185         1  10.10.1.1:1701         1  Established

```

show services l2tp tunnel detail (LAC)

```

user@host> show services l2tp tunnel detail
Tunnel local ID: 31889, Tunnel remote ID:      1
Remote IP: 100.1.1.1:1701
Sessions: 1, State: Established
Tunnel Name: 1/tunnel-to-LNS-1
Local IP: 100.1.1.2:1701
Local name: ce-lac, Remote name: ce-lns
Effective Peer Resync Mechanism: silent failover

```

show services l2tp tunnel detail (LAC on MX Series Routers)

```

user@host> show services l2tp tunnel detail
Tunnel local ID: 17301, Tunnel remote ID: 1
Remote IP: 10.10.1.1:1701
Sessions: 1, State: Established
Tunnel Name: 2/tunnel-to-LNS-2
Local IP: 100.1.1.2:1701
Local name: ce-lac, Remote name: ce-lns
Effective Peer Resync Mechanism: silent failover
Tunnel Logical System: default, Tunnel Routing Instance: default

```

show services l2tp tunnel detail (LNS on MX Series Routers)

```

user@host> show services l2tp tunnel detail
Tunnel local ID: 17301, Tunnel remote ID: 1
Remote IP: 12.1.1.15:1701

```

```
Sessions: 1, State: Established
Tunnel Name: 2/2
Local IP: 12.1.1.5:1701
Local name: ce-bras-mx240-e, Remote name: testlac2
Effective Peer Resync Mechanism: silent failover
Tunnel Logical System: default, Tunnel Routing Instance: vrf1
```

show services l2tp tunnel extensive (LAC)

```
user@host> show services l2tp tunnel extensive
Tunnel local ID: 17185, Tunnel remote ID: 1
Remote IP: 10.10.1.1:1701
Sessions: 1, State: Established
Tunnel Name: 2/tunnel-to-LNS-2
Local IP: 100.1.1.2:1701
Local name: ce-lac, Remote name: ce-lns
Effective Peer Resync Mechanism: failover protocol
Max sessions: 32000, Window size: 4, Hello interval: 60
Create time: Tue Nov 9 15:23:29 2010, Up time: 00:00:26
Idle time: 00:00:00
```

show services l2tp tunnel extensive (LNS on M Series Routers)

```
user@host> show services l2tp tunnel extensive
Interface: sp-1/2/0, Tunnel group: group1
Tunnel local ID: 62746, Tunnel remote ID: 16930
Remote IP: 10.128.1.2:1701
Sessions: 1, State: Established
Local IP: 10.128.1.1:1701
Local name: router-1, Remote name: router-2
Max sessions: 50, Window size: 32, Hello interval: 60
Create time: Tue Mar 23 14:13:15 2004, Up time: 01:14:58
Idle time: 00:00:07
Statistics since: Tue Mar 23 14:13:13 2004
```

| | Packets | Bytes |
|------------|---------|-------|
| Control Tx | 80 | 1152 |
| Control Rx | 3 | 272 |
| Data Tx | 0 | 0 |
| Data Rx | 450 | 28.0k |
| Errors Tx | 0 | |
| Errors Rx | 0 | |

```
Interface: sp-1/2/0, Tunnel group: group_company_dns
Tunnel local ID: 37266, Tunnel remote ID: 36217
Remote IP: 10.128.11.2:1701
Sessions: 1, State: Established
Local IP: 10.128.11.1:1701
Local name: router-1, Remote name: router-2
Max sessions: unlimited, Window size: 32, Hello interval: 60
Create time: Tue Mar 23 14:13:15 2004, Up time: 01:14:59
Idle time: 01:14:55
Statistics since: Tue Mar 23 14:13:13 2004
```

| | Packets | Bytes |
|------------|---------|-------|
| Control Tx | 81 | 1164 |
| Control Rx | 3 | 273 |
| Data Tx | 0 | 0 |
| Data Rx | 1 | 80 |
| Errors Tx | 0 | |
| Errors Rx | 0 | |

show services l2tp tunnel extensive (LNS on MX Series Routers)

```
user@host> show services l2tp tunnel extensive
Tunnel local ID: 40553, Tunnel remote ID: 1
Remote IP: 192.168.1.3:1701
Sessions: 1, State: Established
Tunnel Name: 3/1838
Local IP: 10.1.1.2:1701
Local name: lns-mx960, Remote name: testlac
Effective Peer Resync Mechanism: silent failover
Nas Port Method: none
Tunnel Logical System: default, Tunnel Routing Instance: vrf1
Max sessions: 60000, Window size: 4, Hello interval: 60
Create time: Mon Apr 25 20:27:50 2011, Up time: 00:01:11
Idle time: 00:00:00, ToS Reflect: Enabled
Tunnel Group Name: tg1
Statistics since: Mon Apr 25 20:27:50 2011
```

| | Packets | Bytes |
|------------|---------|-------|
| Control Tx | 4 | 219 |
| Control Rx | 4 | 221 |
| Data Tx | 0 | 0 |
| Data Rx | 6 | 64 |
| Errors Tx | 0 | |
| Errors Rx | | |

show services l2tp tunnel statistics (MX Series Routers)

```
user@host>show services l2tp tunnel statistics
Tunnel local ID: 17185, Tunnel remote ID: 1
Sessions: 31.8k, State: Established
Statistics since: Mon Aug 1 13:21:38 2011
```

| | Packets | Bytes |
|------------|---------|---------|
| Control Tx | 90.3k | 9.0M |
| Control Rx | 32.0k | 1296.9k |
| Data Tx | 127.3k | 1591.6k |
| Data Rx | 100.8k | 1273.4k |
| Errors Tx | 0 | |
| Errors Rx | 0 | |

show services l2tp user

| | |
|---------------------------------|--|
| Syntax | show services l2tp user
<brief detail extensive statistics>
<user <i>username</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M10i and M7i routers only) Display a list of active Layer 2 Tunneling Protocol (L2TP) users. |
| Options | <p>none—Display all active L2TP users.</p> <p>brief detail extensive statistics—(Optional) Display the specified level of output. Use the statistics option to display L2TP user statistics.</p> <p>user <i>username</i>—(Optional) Display L2TP user information for only the specified username.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • L2TP Services Configuration Overview on page 638 • L2TP Minimum Configuration on page 639 |
| List of Sample Output | show services l2tp user extensive on page 2005 |
| Output Fields | Table 87 on page 2003 lists the output fields for the show services l2tp user command. Output fields are listed in the approximate order in which they appear. |

Table 87: show services l2tp user Output Fields

| Field Name | Field Description |
|--------------------------|--|
| Interface | Name of an adaptive services interface. |
| Tunnel group | Name of a tunnel group. |
| Tunnel local ID | Local identifier of the tunnel, as assigned by the L2TP network server (LNS). |
| Session local ID | Local identifier of the session, as assigned by the L2TP network server (LNS). |
| Session remote ID | Remote identifier of the session, as assigned by the L2TP access concentrator (LAC). |

Table 87: show services l2tp user Output Fields (*continued*)

| Field Name | Field Description |
|----------------------------|---|
| State | State of the L2TP session: <ul style="list-style-type: none"> • Established—The session is operating. • closed—The session is being closed. • destroyed—The session is being destroyed. • clean-up—The session is being cleaned up. • Ins-ic-accept-new—A new session is being accepted. • Ins-ic-idle—The session has been created and is idle. • Ins-ic-reject-new—The new session is being rejected. • Ins-ic-wait-connect—The session is waiting for the peer's incoming call connected (ICCN) message. |
| Mode | Mode of the interface representing the session: shared or exclusive . |
| Local IP | IP address of the local endpoint of the tunnel. |
| Remote IP | IP address of the peer endpoint of the tunnel. |
| Username | Name of the user logged in to the session. |
| Assigned IP address | IP address assigned to remote client. |
| Local name | Name of the local device. |
| Remote name | Name of the remote device. |
| Local MRU | Maximum receive unit (MRU) setting of the local device, in bytes. |
| Remote MRU | MRU setting of the remote device, in bytes. |
| Tx speed | Transmit speed of the tunnel session, in bps. |
| Rx speed | Receive speed of the tunnel session, in bps. |
| Bearer type | Type of bearer enabled: <ul style="list-style-type: none"> • 0—Might indicate that the call was not received over a physical link (for example, when the LAC and PPP are located in the same subsystem) • 1—Digital access requested • 2—Analog access requested • 4—Asynchronous Transfer Mode (ATM) bearer support |
| Framing type | Type of framing enabled: <ul style="list-style-type: none"> • 1—Synchronous framing • 2—Asynchronous framing |
| LCP renegotiation | Whether Link Control Protocol (LCP) renegotiation is configured: On or Off . |

Table 87: show services l2tp user Output Fields (*continued*)

| Field Name | Field Description |
|---------------------------|---|
| Authentication | Type of authentication algorithm used: Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP). |
| Interface ID | Name of the logical unit. |
| Interface unit | Logical unit number. |
| Call serial number | Unique serial number assigned to the call. |
| Create time | Date and time when the call was created. |
| Up time | Amount of time elapsed since the call became active, in hours, minutes, and seconds. |
| Idle time | Amount of time elapsed since the call became idle, in hours, minutes, and seconds. |
| Statistics since | <p>Date and time when collection of the following statistics began:</p> <ul style="list-style-type: none"> • Control Tx—Amount of control information transmitted, in packets and bytes. • Control Rx—Amount of control information received, in packets and bytes. • Data Tx—Amount of data transmitted, in packets and bytes. • Data Rx—Amount of data received, in packets and bytes. • Errors Tx—Number of errors transmitted, in packets. • Errors Rx—Number of errors received, in packets. |

Sample Output

show services l2tp user extensive

```

user@host> show services l2tp user extensive
Interface: sp-1/2/0, Tunnel group: group1, Tunnel local ID: 62746
Session local ID: 56793, Session remote ID: 53304
State: Established, Mode: shared
Local IP: 10.128.1.1:1701, Remote IP: 10.128.1.2:1701
Username: usr1@juniper_1.net, Assigned IP address: 10.50.2.1/32
Local name: router-1, Remote name: router-2
Local MRU: 4000, Remote MRU: 1500, Tx speed: 64000, Rx speed: 64000
Bearer type: 2, Framing type: 1
LCP renegotiation: Off, Authentication: CHAP, Interface ID: unit_20
Interface unit: 20, Call serial number: 4137941434
Create time: Tue Mar 23 14:13:15 2004, Up time: 01:16:41
Idle time: 00:00:00
Statistics since: Tue Mar 23 14:13:13 2004

```

| | Packets | Bytes |
|------------|---------|-------|
| Control Tx | 4 | 88 |
| Control Rx | 2 | 28 |
| Data Tx | 0 | 0 |
| Data Rx | 461 | 29.0k |
| Errors Tx | 0 | |
| Errors Rx | 0 | |

```

Interface: sp-1/2/0, Tunnel group: group_company_dns, Tunnel local ID: 37266
Session local ID: 39962, Session remote ID: 53303

```

State: Established, Username: usr1@company_dns.com, Mode: shared
Local IP: 10.128.11.1:1701, Remote IP: 10.128.11.2:1701
Username: usr1@company_dns.com, Assigned IP address: 10.48.1.1/32
Local name: router-1, Remote name: router-2
Local MRU: 4470, Remote MRU: 4470, Tx speed: 155000000,
Rx speed: 155000000
Bearer type: 2, Framing type: 1
LCP renegotiation: Off, Authentication: CHAP, Interface ID: unit_31
Interface unit: 31, Call serial number: 4137941433
Create time: Tue Mar 23 14:13:17 2004, Up time: 01:16:39
Idle time: 01:16:36
Statistics since: Tue Mar 23 14:13:15 2004

| | Packets | Bytes |
|------------|---------|-------|
| Control Tx | 6 | 196 |
| Control Rx | 4 | 150 |
| Data Tx | 0 | 0 |
| Data Rx | 1 | 80 |
| Errors Tx | 0 | |
| Errors Rx | 0 | |

show services nat ipv6-multicast-interfaces

| | |
|---------------------------------|---|
| Syntax | show services nat ipv6-multicast-interfaces |
| Release Information | Command introduced in Junos OS Release 8.5. |
| Description | Displays a list of interfaces enabled for IPv6 mutlicast. |
| Required Privilege Level | view |
| List of Sample Output | show services nat ipv6-multicast-interfaces on page 2007 |
| Output Fields | Table 88 on page 2007 lists the output fields for the show services nat ipv6-multicast-interfaces command. Output fields are listed in the approximate order in which they appear. |

Table 88: show services nat ipv6-multicast-interfaces Output Fields

| Field Name | Field Description | Level of Output |
|--------------------------|--|-----------------|
| Interface | Name of a service interface. | All levels |
| Admin State | Configured IPv6 multicast capability of an interface , | All levels |
| Operational State | Operation IPv6 multicast status of an interface. | All levels |

Sample Output

show services nat ipv6-multicast-interfaces

```

user@host> show services nat ipv6-multicast-interfaces
Interface           Admin      Operational
                   State      State
ge-5/1/9            Enabled    Enabled
ge-5/1/8            Enabled    Enabled
ge-5/1/7            Enabled    Enabled
ge-5/1/6            Enabled    Enabled
ge-5/1/5            Enabled    Enabled
ge-5/1/4            Enabled    Enabled
ge-5/1/3            Enabled    Enabled
ge-5/1/2            Enabled    Enabled
ge-5/1/1            Enabled    Enabled
ge-5/1/0            Enabled    Enabled
ge-5/0/9            Enabled    Enabled
ge-5/0/8            Enabled    Enabled
ge-5/0/7            Enabled    Enabled
ge-5/0/6            Enabled    Enabled
ge-5/0/5            Enabled    Enabled
ge-5/0/4            Enabled    Enabled
ge-5/0/3            Enabled    Enabled
ge-5/0/2            Enabled    Enabled
ge-5/0/1            Enabled    Enabled
ge-5/0/0            Enabled    Enabled
ge-1/3/9            Enabled    Enabled

```

| | | |
|----------|---------|---------|
| ge-1/3/8 | Enabled | Enabled |
| ge-1/3/7 | Enabled | Enabled |
| ge-1/3/6 | Enabled | Enabled |
| ge-1/3/5 | Enabled | Enabled |
| ge-1/3/4 | Enabled | Enabled |
| ge-1/3/3 | Enabled | Enabled |
| ge-1/3/2 | Enabled | Enabled |
| ge-1/3/1 | Enabled | Enabled |
| ge-1/3/0 | Enabled | Enabled |
| ge-1/2/9 | Enabled | Enabled |
| ge-1/2/8 | Enabled | Enabled |
| ge-1/2/7 | Enabled | Enabled |
| ge-1/2/6 | Enabled | Enabled |
| ge-1/2/5 | Enabled | Enabled |
| ge-1/2/4 | Enabled | Enabled |
| ge-1/2/3 | Enabled | Enabled |
| ge-1/2/2 | Enabled | Enabled |
| ge-1/2/1 | Enabled | Enabled |
| ge-1/2/0 | Enabled | Enabled |
| ge-1/1/9 | Enabled | Enabled |
| ge-1/1/8 | Enabled | Enabled |
| ge-1/1/7 | Enabled | Enabled |
| ge-1/1/6 | Enabled | Enabled |
| ge-1/1/5 | Enabled | Enabled |
| ge-1/1/4 | Enabled | Enabled |
| ge-1/1/3 | Enabled | Enabled |
| ge-1/1/2 | Enabled | Enabled |
| ge-1/1/1 | Enabled | Enabled |
| ge-1/1/0 | Enabled | Enabled |
| ge-1/0/9 | Enabled | Enabled |
| ge-1/0/8 | Enabled | Enabled |
| ge-1/0/7 | Enabled | Enabled |
| ge-1/0/6 | Enabled | Enabled |
| ge-1/0/5 | Enabled | Enabled |
| ge-1/0/4 | Enabled | Enabled |
| ge-1/0/3 | Enabled | Enabled |
| ge-1/0/2 | Enabled | Enabled |
| ge-1/0/1 | Enabled | Enabled |
| ge-1/0/0 | Enabled | Enabled |
| xe-0/3/0 | Enabled | Enabled |
| xe-0/2/0 | Enabled | Enabled |
| xe-0/1/0 | Enabled | Enabled |
| xe-0/0/0 | Enabled | Enabled |

show services nat mappings

| | |
|----------------------------|---|
| Syntax | <pre>show services nat mappings <brief detail summary> <nptv6 (ipv6-address external ipv6-address internal ipv6-address)> <pool-name> <address-pooling-paired endpoint-independent pcp></pre> |
| Release Information | <p>Command introduced in Junos OS Release 10.1.</p> <p>summary option introduced in Junos OS Release 11.1.</p> <p>address-pooling paired option introduced in Junos OS Release 13.2.</p> <p>endpoint-independent option introduced in Junos OS Release 13.2.</p> <p>pcp option introduced in Junos OS Release 13.2.</p> <p>nptv6 option introduced in Junos OS Release 15.1.</p> |
| Description | Display information about Network Address Translation (NAT) address, port, and port control protocol (PCP) mappings. |
| Options | <p>none—Display standard information about all NAT pools.</p> <p>brief detail summary—(Optional) Display the specified level of output.</p> <p>nptv6—(Optional) Display information about the network prefix translation for IPv6 traffic.</p> <p>ipv6-address—(Optional) Display the network prefix translation details for the specified IPv6 address.</p> <p>external—(Optional) Display the external to internal address mapping for a given external address if the mapping exists for stateless network IPv6 prefix translation.</p> <p>internal—(Optional) Display the internal to external address mapping for a given internal address if the mapping exists for stateless network IPv6 prefix translation.</p> <p>pool-name—(Optional) Display detailed information about a specific NAT pool. Used only with detail level output.</p> <p>address-pooling-paired—(Optional) Display only information about address-pooling paired mappings.</p> <p>endpoint-independent—(Optional) Display only information about endpoint-independent mappings.</p> <p>pcp—(Optional) Display only information about port control protocol mappings.</p> |



NOTE: PCP requests with the prefer-failure option request a particular external IP address and port. When the request cannot be fulfilled, the mapping is not created. In this case, the subscriber does not have a mapped IP address. Such a subscriber is counted in the summary of the number or address mappings, but is not displayed in the list of address mappings, as shown in the following examples:

```
user@host# show services nat mappings summary
Service Interface:                               sp-2/0/0
Total number of address mappings:                 1
Total number of endpoint independent port mappings: 0
Total number of endpoint independent filters:      0

user@host# show services nat mappings address-pooling-paired
[edit]
```

This is expected behavior because unfulfilled address mappings (IP of 0.0.0.0) are not displayed in the output of the second CLI command. These address mappings will time out based on configured or default values.

Required Privilege Level view

List of Sample Output

- [show services nat mappings brief on page 2011](#)
- [show services nat mapping detail on page 2012](#)
- [show services nat mappings pool-name on page 2012](#)
- [show services nat mappings summary on page 2012](#)
- [show services nat mappings address-pooling-paired on page 2012](#)
- [show services nat mappings address-pooling-paired \(mapping of active B4 for a subscriber\) on page 2012](#)
- [show services nat mappings endpoint-independent on page 2013](#)
- [show services nat mappings pcp on page 2013](#)
- [show services nat mappings nptv6 internal on page 2013](#)
- [show services nat mappings nptv6 external on page 2013](#)

Output Fields [Table 89 on page 2010](#) lists the output fields for the **show services nat mappings** command. Output fields are listed in the approximate order in which they appear.

Table 89: show services nat mappings Output Fields

| Field Name | Field Description | Level of Output |
|-------------|---|-----------------|
| Interface | Name of a service interface. | All levels |
| Service set | Name of a service set. Individual empty service sets are not displayed, but if none of the service sets has any flows, a flow table header is printed for each service set. | All levels |
| NAT pool | Name of the NAT pool. | All levels |

Table 89: show services nat mappings Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|--|-----------------|
| Address Mapping
or
Mapping | Mapping performed by NAT to conceal the network address. | All levels |
| No. of port mappings | Number of port mappings. | All levels |
| Port mapping | Port mapping performed by NAT. | detail |
| Flow Count | Number of flows. | detail |
| Total number of address mappings | Total number of address mappings, by service interface. | summary |
| Total number of endpoint independent port mappings: | Total number of port mappings by service interface. | summary |
| Total number of endpoint independent filters | Total number of independent filters that filter out only packets that are not destined to the internal address and port, regardless of the external IP address and port source, by service interface. | summary |
| Mapping State | NAT mapping state. The following states are possible: <ul style="list-style-type: none"> • ACTIVE—Indicates that the entry is active and in use. • TIMEOUT—Indicates that the mapping is not in use. After the mapping-timeout, configured at the [edit services nat pool pool-name] hierarchy level, lapses, the mapping is deleted. This field also displays the number of seconds after which the timeout occurs. | |
| Ports In Use | The number of ports used for a specific address-pooling paired mapping. | |
| PCP Lifetime | Elapsed PCP lifetime in seconds. | |
| PCP Client | Address of the PCP client sending the PCP request. | |
| Session Count | Number of sessions currently using the mapping. | |

Sample Output

show services nat mappings brief

```

user@host> show services nat mappings brief
Interface: sp-2/3/0, Service set: s1

NAT pool: p1
  Address Mapping: 2.1.20.10 ---> 34.34.34.34
  No. of port mappings: 1

```

show services nat mapping detail

```
user@host> show services nat mapping detail
Interface: sp-2/3/0, Service set: s1

NAT pool: p1
Address Mapping: 2.1.20.10 ---> 34.34.34.34, No. of port mappings: 1
Port mapping: 49604 --> 1024, Flow Count: 2
```

show services nat mappings pool-name

```
user@host> show services nat mappings pool-name p1
Interface: sp-2/3/0, Service set: s1

NAT pool: p1
Address Mapping: 2.1.20.10 ---> 34.34.34.34
No. of port mappings: 1
```

show services nat mappings summary

```
user@host> show services nat mapping summary

Service Interface:                                sp-1/0/0
Total number of address mappings:                  790
Total number of endpoint independent port mappings: 1580
Total number of endpoint independent filters:       1580

Service Interface:                                sp-1/1/0
Total number of address mappings:                  914
Total number of endpoint independent port mappings: 1828
Total number of endpoint independent filters:       1828

Service Interface:                                sp-4/0/0
Total number of address mappings:                  688
Total number of endpoint independent port mappings: 1376
Total number of endpoint independent filters:       1376

Service Interface:                                sp-4/1/0
Total number of address mappings:                  648
Total number of endpoint independent port mappings: 1296
Total number of endpoint independent filters:       1296
```

show services nat mappings address-pooling-paired

```
user@host> show services nat mappings address-pooling-paired
Interface: sp-3/0/0, Service set: NAPT44-SS1
NAT pool: napt44-SS1-p1
Mapping      : 29.32.38.255    --> 192.168.75.23
Ports In Use :      9
Session Count:      1
Mapping State: Active
```

show services nat mappings address-pooling-paired (mapping of active B4 for a subscriber)

```
user@host> show services nat mappings address-pooling-paired
Interface: sp-0/0/0, Service set: sset_1

NAT pool: nat_pool1

Mapping      : 2001::          --> 33.33.33.2
```

```

Ports In Use      :      1
Session Count     :      9
Mapping State     : Timeout

```

show services nat mappings endpoint-independent

```

user@host> show services nat mappings endpoint-independent
Interface: sp-3/0/0, Service set: NAPT44-SS1
NAT pool: napt44-SS1-p1
Mapping          : 29.32.38.255:10000    --> 192.168.75.23:1024
Session Count    : 1
Mapping State    : Active

```

show services nat mappings pcsp

```

user@host> show services nat mappings pcsp
PCP Client       : 172.16.0.1           PCP Lifetime : 45
Mapping          : 29.32.38.255:10000    --> 192.168.75.23:1024
Session Count    : 1
Mapping State    : Active

```

show services nat mappings nptv6 internal

```

user@host> show services nat mappings nptv6 internal 1111:2222:3333:aaaa:bbbb::1

Interface      Service-set  NAT-Pool      Address Mapping
si-0/1/0       ss_nptv6    ss_nptv6_pool 1111:2222:3333:aaaa:bbbb::1 ->
aaaa:bbbb:cccc:dddd:bbbb::1

```

show services nat mappings nptv6 external

```

user@host> show services nat mappings nptv6 external aaaa:bbbb:cccc:dddd:bbbb::1

Interface      Service-set  NAT-Pool      Address Mapping
si-0/1/0       ss_nptv6    ss_nptv6_pool 1111:2222:3333:aaaa:bbbb::1
-> aaaa:bbbb:cccc:dddd:bbbb::1

```

show services nat pool

| | |
|----------------------------|---|
| Syntax | <code>show services nat pool</code>
<code><brief detail></code>
<code><pool-name></code>
<code>pgcp <ports-per-session remotely-controlled></code> |
| Release Information | Command introduced before Junos OS Release 7.4.
<code>pgcp</code> option added in Junos OS Release 8.5. |
| Description | Display information about Network Address Translation (NAT) pools. |



NOTE: On MS-MPCs and MS-MICs, if the line cards receive a packet immediately after the active port block timeout interval has expired, a new port block is allocated and the old port block is released thereafter (if no more ports are being used from that block). In such a scenario, you might notice that the Max number of port blocks used field displays a higher value than the value shown for the Unique pool users field in the output of the `show services nat pool detail` command. This behavior is expected with port block allocation.

With MS-MPCs and MS-MICs, in the output of the `show services nat pool detail` command, the Max ports used and the Ports in use fields display values that indicate a higher number than the number of active subscribers on the member interfaces of an ams interface. This behavior of an increased value displayed for the number of ports allocated and maximum number of ports used is expected after you perform a Graceful Routing Engine switchover (GRES) and a restart of the MPC.

With MS-MPCs and MS-MICs on MX Series routers with AMS interfaces, it is observed that the subscriber and port count details are displayed only after a long time in the output of the `show services nat pool detail` command. This behavior is expected with NAT pool counters and occurs, regardless of port block allocation being configured.

| | |
|----------------|---|
| Options | none —Display standard information about all NAT pools. |
| | brief detail —(Optional) Display the specified level of output. |
| | pool-name —(Optional) Display information about the specified NAT pool. |
| | pgcp —(Optional) Display information about a NAT pool that is exclusive to the BGF. |
| | ports-per-session —(Optional) Display the number of ports allocated per session from the NAT pool. |
| | remotely-controlled —(Optional) Display if the NAT pool is explicitly specified by the gateway controller. |

Required Privilege Level view

List of Sample Output [show services nat pool brief on page 2016](#)
[show services nat pool detail on page 2017](#)
[show services nat pool for Secured Port Block Allocation on page 2017](#)
[show services nat pool detail for Deterministic Port Block Allocation on page 2017](#)
[show services nat pool for Deterministic Port Block Allocation on page 2018](#)
[show services nat pool detail for Port Block Allocation on page 2018](#)

Output Fields [Table 90 on page 2015](#) lists the output fields for the **show services nat pool** command. Output fields are listed in the approximate order in which they appear.

Table 90: show services nat pool Output Fields

| Field Name | Field Description | Level of Output |
|---|--|-----------------|
| DetNat subscriber exceeded port limits | The number of times a subscriber exceeded its port limits for a NAT pool that uses deterministic port block allocation. | All levels. |
| MAX number of port blocks used | The maximum number of port blocks used. | All levels. |
| Port block memory allocation errors | The number of memory allocation errors for port blocks. | All levels. |
| Current number of port blocks in use | Current count of the port blocks that are being used. | |
| Unique pool users | The number of different users of the NAT pools. | All levels. |
| Port block allocation errors | The consolidated number of port block allocation errors. | All levels. |
| Port blocks limit exceeded errors | The total number of times when a request for more than the allowed port blocks allocated for a user arrives from a user. | All levels. |
| Interface | Name of an adaptive services interface. | All levels |
| Service set | Name of a service set. Individual empty service sets are not displayed, but if none of the service sets has any flows, a flow table header is printed for each service set. | All levels |
| NAT pool | Name of the Network Address Translation pool. | All levels |
| Type or Translation type | Address translation type: basic-nat-pt , basic-nat44 , basic-nat66 , deterministic-napt44 , dnat-44 , dynamic-nat44 , napt44 , napt-66 , napt-pt , stateful-nat64 , twice-basic-nat-44 , twice-dynamic-nat-44 , twice-dynamic-napt-44 . | All levels |
| Address or Address range | IPv4 address range of the pool. | All levels |

Table 90: show services nat pool Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|---|-----------------|
| Configured port range | The range of ports configured to be used for NAT pool. | detail |
| Preserve range enabled | Whether the capability to preserve the privileged port range after translation is enabled. One of the following is displayed: <ul style="list-style-type: none"> • Is active—Preservation of port range is enabled. • Not active—Preservation of port range is not enabled. | detail |
| Port or Port range | Port range of the pool. Applicable only for dynamic NAT pools. Not displayed for static NAT pools. | All levels |
| Ports used' or Ports in use | Number of ports allocated in this pool with this name. Applicable only for dynamic NAT pools. Not displayed for static NAT pools. | All levels |
| Port block type | Type of port block allocation: secured or deterministic | All levels |
| Out of port errors | Number of port allocation errors. Applicable only for dynamic NAT pools. Not displayed for static NAT pools. | detail |
| Max ports used | Maximum number of ports used. Applicable only for dynamic NAT pools. Not displayed for static NAT pools. | detail |
| Addresses in use | Number of addresses in use for dynamic source address NAT pools. | detail |
| Out of Port Errors | No more ports available to allocate. | Detail |
| Max Ports Used | The maximum number of ports in use at any time since the services PIC was started. | Detail |
| AP-P out of port errors | When address pooling paired (AP-P) is configured, a private IP is paired to a public IP. This is counter of translation errors where there are free ports available in the NAT pool, but none for the NAT IP to which the private IP is paired. | Detail |
| Current EIF Inbound flows count | Current count of EIF inbound flows, including all EIF flows per pool. | |
| EIF flow limit exceeded drops | Current number of flow drops due to exceeded flow limit. This number is per pool, not per EIF mapping. | |

Sample Output

show services nat pool brief

```
user@host> show services nat pool brief
```

```
Interface: ms-1/0/0, Service set: s1
NAT pool      Type   Address          Port      Ports used
dest-pool     DNAT-44 10.10.10.2-10.10.10.2
napt-pool     NAPT-44 50.50.50.1-50.50.50.254 1024-63487 0
```

```
source-dynamic-pool DYNAMIC NAT44 40.40.40.1-40.40.40.254
source-static-pool BASIC NAT44 30.30.30.1-30.30.30.254
```

show services nat pool detail

```
user@host> show services nat pool detail

Interface: ms-1/0/0, Service set: s1
  NAT pool: dest-pool, Translation type: DNAT-44
    Address range: 10.10.10.2-10.10.10.2
    Configured port range: 1-60000, Preserve range enabled: Is active
  NAT pool: napt-pool, Translation type: NAPT-44
    Address range: 50.50.50.1-50.50.50.254
    Configured port range: 1-60000, Preserve range enabled: Is active
    Port range: 1024-63487, Ports in use: 0, Out of port errors: 0, Max ports
used: 0
  NAT pool: source-dynamic-pool, Translation type: DYNAMIC NAT44
    Address range: 40.40.40.1-40.40.40.254
    Configured port range: 1-60000, Preserve range enabled: Is active
    Out of address errors: 0, Addresses in use: 0
  NAT pool: source-static-pool, Translation type: BASIC NAT44
    Address range: 30.30.30.1-30.30.30.254
    Configured port range: 1-60000, Preserve range enabled: Is active
```

show services nat pool for Secured Port Block Allocation

```
user@host> show services nat pool

Interface: sp-2/0/0, Service set: in
  NAT pool      Type      Address      Port      Ports used
  mypool        dynamic  3.3.3.3-3.3.3.10  512-65535  0
                3.3.3.15-3.3.3.20
                3.3.3.25-3.3.3.30
                3.3.3.95-3.3.3.200
Port block size: 64, Max port blocks per address: 1, Active block timeout: 86400,
Effective port range: 1024-65471,
Effective number of port blocks: 126882, Effective number of ports: 8120448, Port
block efficiency: nan

Interface: sp-2/1/0, Service set: in1
  NAT pool      Type      Address      Port      Ports used
  mypool1       dynamic  9.9.9.1-9.9.9.254  512-65535  0
Port block size: 64, Max port blocks per address: 1, Active block timeout: 86400,
Effective port range: 1024-65471,
Effective number of port blocks: 255778, Effective number of ports: 16369792,
Port block efficiency: nan
```

show services nat pool detail for Deterministic Port Block Allocation

```
user@host> show services nat pool detail

Interface: sp-2/0/0, Service set: ss1
  NAT pool: napt_pool, Translation type: dynamic
    Address range: 5.5.5.1-5.5.5.254
    Configured port range: 1-60000, Preserve range enabled: Is active
    Port range: 2000-2002, Ports in use: 2, Out of port errors: 0, Max ports used:
2
    AP-P out of port errors: 188
    Max number of port blocks used: 1, Current number of port blocks in use: 1,
Port block allocation errors: 0,
Port block memory allocation errors: 0
DetNAT subscriber exceeded port limits: 1 <<<<<<<<<
Unique pool users: 1
```

show services nat pool for Deterministic Port Block Allocation

```
user@host> show services nat pool
```

```
Interface: sp-2/0/0, Service set: ss2
NAT pool      Type      Address                      Port      Ports Used
pba           dynamic  33.33.33.1-33.33.33.128    512-65535 6604
Port block type: Deterministic port block, Port block size: 200
```

show services nat pool detail for Port Block Allocation

```
user@host> show services nat pool detail
```

```
Interface: sp-2/0/0, Service set: s
NAT pool: napt_pool, Translation type: dynamic
Address range: 44.1.1.1-44.1.1.1
Configured port range: 1-60000
Port range: 1024-65535, Ports in use: 0, Out of port errors: 0,
Max ports used: 0
AP-P out of port errors: 0
Current EIF Inbound flows count: 0
EIF flow limit exceeded drops: 0
```

Sample Output

show services pcsp statistics

| | |
|---------------------------------|---|
| Syntax | show services pcsp statistics |
| Release Information | Command introduced in Junos OS Release 13.2 |
| Description | Display information PCP mappings. |
| Required Privilege Level | view |
| List of Sample Output | show services pcsp statistics pcsp on page 2020 |
| Output Fields | Table 91 on page 2019 lists the output fields for the show services pcsp statistics command. Output fields are listed in the approximate order in which they appear. |

Table 91: show services pcsp statistics Output Fields

| Field Name | Field Description |
|-------------------------------|--|
| Services PIC Name | Name of a service interface. |
| Protocol Statistics | Overall PCP statistics, consisting of: operational, option, and results statistics. |
| Operational Statistics | Operational statistics group. |
| Map request received | Total PCP MAP requests received from PCP clients. |
| Peer request received | Number of peer requests received. |
| Option Statistics | Number of requests using available options. |
| Unprocessed requests received | Number of requests received with no option specified. |
| Third party requests received | Number of third-party requests received. |
| Prefer fail option received | Number of prefer fail requests received. |
| Filter option received | Number of filter option requests received. |
| Other options counters | Number of packets received with options other than prefer-fail and third-party . |
| Other optional received | |
| Results Statistics | Information about the results of PCP requests. |
| PCP success | Number of PCP MAP requests successfully processed by the server. |
| PCP unsupported version | Number of PCP packets received with version other than 1. |
| Not authorized | Number of unauthorized MAP delete requests. |

Table 91: show services pcp statistics Output Fields (*continued*)

| Field Name | Field Description |
|----------------------------------|--|
| Bad requests | Number of requests with invalid PCP packets. |
| Unsupported opcode | Number of packets that have an unsupported opcode. |
| Unsupported option | Number of packets that have an unsupported option. |
| Bad option | Number of packet that have a malformed option. |
| Network failure | Number of times a mapping could not be provided due to a network failure. |
| Out of resources | Number of times a mapping could not be provided because the PCP server ran out of pool resources. |
| Unsupported protocol | Number of requests for which the protocol was neither TCP nor UDP. |
| User exceeded quota | Number of requests for which the PCP client requested more than the configured number of ports. |
| Cannot provide external | Number of requests for which the PCP server cannot provide the external address or port requested by the client. |
| Address mismatch | Number of requests for which the PCP client IP address and the layer-3 source IP do not match. |
| Excessive number of remote peers | This counter is not currently used. |
| Processing error | Number of requests with malformed PCP packets information, such as an invalid IP address in a third-party request . |
| Other result counters | Not currently used. |

Sample Output

show services pcp statistics pcp

```
user@host> show services pcp statistics pcp
Services PIC Name:    sp-2/1/0
```

```
Protocol Statistics:
```

```
Operational Statistics
```

```
Map request received           : 0
Peer request received          : 0
Other operational counters     : 0
```

```
Option Statistics
```

```
Unprocessed requests received  : 0
Third party requests received   : 0
```

| | |
|-----------------------------|-----|
| Prefer fail option received | : 0 |
| Filter option received | : 0 |
| Other options counters | : 0 |
| Option optional received | : 0 |

Result Statistics

| | |
|----------------------------------|-----|
| PCP success | : 0 |
| PCP unsupported version | : 0 |
| Not authorized | : 0 |
| Bad requests | : 0 |
| Unsupported opcode | : 0 |
| Unsupported option | : 0 |
| Bad option | : 0 |
| Network failure | : 0 |
| Out of resources | : 0 |
| Unsupported protocol | : 0 |
| User exceeded quota | : 0 |
| Cannot provide external | : 0 |
| Address mismatch | : 0 |
| Excessive number of remote peers | : 0 |
| Processing error | : 0 |
| Other result counters | : 0 |

show services service-sets cpu-usage

| | |
|---------------------------------|--|
| Syntax | show services service-sets cpu-usage
<interface <i>interface-name</i> >
<service-set <i>service-set-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display service set CPU usage as a percentage. The command is supported only on Adaptive Services PICs (SP PICs). |
| Options | <p>none—Display CPU usage for all adaptive services interfaces and service sets.</p> <p>interface <i>interface-name</i>—(Optional) Display CPU usage for a particular interface. On M Series and T Series routers, the <i>interface-name</i> parameter can have the value <i>sp-fpc/pic/port</i> or <i>rspnumber</i>.</p> <p>service-set <i>service-set-name</i>—(Optional) Display CPU usage for a particular service set. For the Layer 2 Tunneling Protocol (L2TP), you can use a tunnel group to represent a service set.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services service-sets cpu-usage on page 2022 |
| Output Fields | Table 92 on page 2022 lists the output fields for the show services service-sets cpu-usage command. Output fields are listed in the approximate order in which they appear. |

Table 92: show services service-sets cpu-usage Output Fields

| Field Name | Field Description |
|-------------------------------|--|
| Interface | Name of an adaptive services interface |
| Service set (system category) | Name of the CPU usage category: <ul style="list-style-type: none"> • idp_recommended—Name of the service sets (displays all the service sets attached to the service PICs) • Idle • System • Receive • Transmit |
| CPU utilization % | Percentage of the CPU resources being used |

Sample Output

show services service-sets cpu-usage

```
user@host> show services service-sets cpu-usage
```


| Interface | Service set (system category) | CPU utilization % |
|-----------|-------------------------------|-------------------|
| sp-4/1/0 | idp_recommended | 18.20 % |
| sp-4/1/0 | Idle | 44.69 % |
| sp-4/1/0 | System | 7.01 % |
| sp-4/1/0 | Receive | 15.10 % |
| sp-4/1/0 | Transmit | 15.00 % |

show services service-sets memory-usage

Syntax show services service-sets memory-usage
 <interface *interface-name*>
 <service-set *service-set-name*>
 <zone>

Release Information Command introduced before Junos OS Release 7.4.

Description Display service set memory usage.

Options none—Display service set memory usage.

interface *interface-name*—(Optional) Display memory usage for a particular interface. On M Series and T Series routers, the *interface-name* can be *sp-fpc/pic/port*, or *rspnumber*.



NOTE: This command is not supported on Multilink Protocol–based services PICs.

The interface option is not supported on Multiservice PICs.

service-set *service-set-name*—(Optional) Display memory usage for a particular service set. For Layer 2 Tunneling Protocol (L2TP), you can use a tunnel group to represent a service set.

zone—(Optional) Display the memory usage zone of the adaptive services interface or an individual service set.

Required Privilege Level view

List of Sample Output [show services service-sets memory-usage on page 2025](#)
[show services service-sets memory-usage zone on page 2025](#)
[show services service-sets memory-usage interface on page 2025](#)

Output Fields [Table 93 on page 2024](#) lists the output fields for the **show services service-sets memory-usage** command. Output fields are listed in the approximate order in which they appear.

Table 93: show services service-sets memory-usage Output Fields

| Field Name | Field Description |
|-------------|--|
| Interface | Name of an adaptive services interface |
| Service set | Name of a service set |
| Bytes Used | Number of bytes of memory being used |

Table 93: show services service-sets memory-usage Output Fields (continued)

| Field Name | Field Description |
|--------------------|--|
| Memory zone | <p>Memory zone in which the adaptive services interface is currently operating:</p> <ul style="list-style-type: none"> • Green—All new flows are allowed. • Yellow—Unused memory is reclaimed. All new flows are allowed. • Orange—New flows are allowed only for service sets that are using less than their equal share of memory. • Red—No new flows are allowed. |

Sample Output

show services service-sets memory-usage

```

user@host> show services service-sets memory-usage
Interface  Service set      Bytes Used
ms-4/0/0   N/A              14817036
ms-4/1/0   N/A              14691700

```

show services service-sets memory-usage zone

```

user@host> show services service-sets memory-usage zone
Interface  Memory zone

```

show services service-sets memory-usage interface

```

user@host> show services service-sets memory-usage interface ms-4/1/0
Interface  Service Set      Bytes Used
ms-4/1/0   N/A              14691700

```

show services service-sets statistics packet-drops

| | |
|---------------------------------|---|
| Syntax | show services service-sets statistics packet-drops
<interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 7.4. |
| Description | Display the number of dropped packets for service sets exceeding CPU limits or memory limits. |
| Options | <p>none—Display the number of dropped service sets packets for all adaptive services interfaces.</p> <p>interface <i>interface-name</i>—(Optional) Display the number of dropped service sets packets for a particular interface. On M Series and T Series routers, <i>interface-name</i> can be <i>ms-fpc/pic/port</i>, <i>sp-fpc/pic/port</i>, or <i>rspnumber</i>.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear services flow-collector statistics on page 2167 |
| List of Sample Output | show services service-sets statistics packet-drops interface on page 2026 |
| Output Fields | Table 94 on page 2026 lists the output fields for the show services service-sets packet-drops command. Output fields are listed in the approximate order in which they appear. |

Table 94: show services service-sets packet-drops Output Fields

| Field Name | Field Description |
|---------------------------|---|
| <i>Interface</i> | Name of an adaptive services interface. |
| <i>Service set</i> | Name of a service set. |
| <i>CPU limit Drops</i> | Number of packets dropped because the service set exceeded the average CPU limit. |
| <i>Memory limit Drops</i> | Number of packets dropped because the service set exceeded the memory limit. |
| <i>Flow limit Drops</i> | Number of packets dropped because the service set exceeded the flow limit. |

Sample Output

show services service-sets statistics packet-drops interface

```
user@host> show services service-sets statistics packet-drops interface sp-1/0/0
```

| Interface | Service Set | Cpu limit
Drops | Memory limit
Drops | Flow limit
Drops |
|-----------|-------------|--------------------|-----------------------|---------------------|
| sp-1/0/0 | sset1 | 0 | 0 | 0 |

show services service-sets statistics syslog

| | |
|---------------------------------|---|
| Syntax | show services service-sets statistics syslog
<interface <i>interface-name</i> >
<service-set <i>service-set-name</i> >
<brief detail> |
| Release Information | Command introduced in Junos OS Release 11.1. |
| Description | Display the system log statistics with optional filtering by interface and service set name.. |
| Options | <p>none—Display the system log statistics for all services interfaces and all service sets.</p> <p>brief—(Default) Display abbreviated system log statistics.</p> <p>detail—Display detailed system log statistics.</p> <p>interface <i>interface-name</i>—(Optional) Display the system log statistics for a specific adaptive service interface. On M Series and T Series routers, <i>interface-name</i> can be <i>ms-fpc/pic/port</i>, <i>sp-fpc/pic/port</i>, or <i>rspnumber</i>.</p> <p>service-set <i>service-set name</i>—(Optional) Display the system log statistics for a specific named service-set.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear services service-sets statistics syslog on page 1853 |
| List of Sample Output | show services service-sets statistics syslog brief on page 2029
show services service-sets statistics syslog detail on page 2029 |
| Output Fields | Table 95 on page 2028 lists the output fields for the show services service-sets statistics syslog command. Output fields are listed in the approximate order in which they appear. |

Table 95: show services service-sets statistics syslog Output Fields

| Field Name | Field Description | Level |
|--------------------|--|-------|
| Interface | Name of a services interface. | all |
| Message rate limit | Maximum number of messages per second written to the interface's system log. | all |
| Service set | Name of a service set. | all |
| Messages sent | Number of messages sent. | brief |
| Messages dropped | Number of messages dropped. | brief |

Table 95: show services service-sets statistics syslog Output Fields (*continued*)

| Field Name | Field Description | Level |
|-------------------|--|--------|
| <i>class name</i> | <p>Logs created for events for each of the following classes:</p> <ul style="list-style-type: none"> • Session open logs • Session close logs • Packet logs • Stateful firewall logs • ALG logs • NAT logs • IDS logs • All other logs <p>The following information is displayed for system log messages for each class of event that is logged:</p> <ul style="list-style-type: none"> • Messages sent—Number of messages sent for session open events. • Messages dropped—Number of messages dropped for session open events. Counts are given for these drop reasons: <ul style="list-style-type: none"> • low priority—The priority of the message was too low for the message to be sent. • no class set—Specific classes of event messages were configured and this class was not selected. • above rate limit—The maximum number of system log messages per second was exceeded. | detail |

Sample Output

show services service-sets statistics syslog brief

```

user@host> show services service-sets statistics syslog brief
Interface: sp-1/1/0
  Message rate limit: 200000
  Service-set: sset-sfw-sp1
    Messages sent: 20
    Messages dropped: 3488
  Service-set: sset-nat-sp1
    Messages sent: 18
    Messages dropped: 91
Interface: sp-1/2/0
  Message rate limit: 15000
  Service-set: sset-sfw-sp2
    Messages sent: 210
    Messages dropped: 579

```

Sample Output

show services service-sets statistics syslog detail

```

user@host> show services service-sets statistics syslog detail
Interface: sp-1/2/0
  Message rate limit: 10
  Service-set: sset-sfw

```

Messages sent: 0
Messages dropped: 1600
Session open logs:
 Sent: 0
 Dropped: 1277 (low priority: 1277, no class set: 0, above rate limit: 0)
Session close logs:
 Sent: 0
 Dropped: 0 (low priority: 0, no class set: 0, above rate limit: 0)
Packet logs:
 Sent: 0
 Dropped: 323 (low priority: 323, no class set: 0, above rate limit: 0)
Stateful firewall logs:
 Sent: 0
 Dropped: 0 (low priority: 0, no class set: 0, above rate limit: 0)
ALG logs:
 Sent: 0
 Dropped: 0 (low priority: 0, no class set: 0, above rate limit: 0)
NAT logs:
 Sent: 0
 Dropped: 0 (low priority: 0, no class set: 0, above rate limit: 0)
IDS logs:
 Sent: 0
 Dropped: 0 (low priority: 0, no class set: 0, above rate limit: 0)
Other logs:
 Sent: 0
 Dropped: 0 (low priority: 0, no class set: 0, above rate limit: 0)

show services service-sets statistics tcp-mss

| | |
|---------------------------------|---|
| Syntax | show services service-sets statistics tcp-mss
<interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 9.5. |
| Description | (M Series and T Series routers only) Display TCP maximum segment size (MSS) statistics for service sets. |
| Options | <p>none—Display service set TCP MSS information for all adaptive services interfaces.</p> <p>interface <i>interface-name</i>—(Optional) Display TCP MSS statistics for a particular interface. The <i>interface-name</i> can be <i>ms-fpc/pic/port</i>, <i>sp-fpc/pic/port</i>, or <i>rsp number</i>.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services service-sets statistics tcp-mss on page 2031 |
| Output Fields | Table 96 on page 2031 lists the output fields for the show services service-sets statistics tcp-mss command. Output fields are listed in the approximate order in which they appear. |

Table 96: show services service-sets statistics tcp-mss Output Fields

| Field Name | Field Description |
|---------------------|--|
| Interface | Name of the adaptive services interface. |
| Service Set | Name of the configured service set. |
| SYN Received | Number of TCP SYN packets received. |
| SYN Modified | Number of TCP SYN packets with the MSS value modified to match the MSS value specified in the TCP MSS configuration. |

Sample Output

show services service-sets statistics tcp-mss

```

user@host> show services service-sets statistics tcp-mss
Interface  Service Set          SYN Received  SYN Modified
sp-1/2/0   asq_ipsec_svc_0      500           220

```

show services service-sets summary

| | |
|---------------------------------|---|
| Syntax | show services service-sets summary
<interface <i>interface-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display service set summary information. |
| Options | <p>none—Display service set summary information for all adaptive services interfaces.</p> <p>interface <i>interface-name</i>—(Optional) Display service set summary information for a particular interface. On M Series and T Series routers, <i>interface-name</i> can be <i>ms-fpc/pic/port</i>, <i>sp-fpc/pic/port</i>, or <i>rspnumber</i>.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services service-sets summary on page 2032
show services service-sets summary interface on page 2033 |
| Output Fields | Table 97 on page 2032 lists the output fields for the show services service-sets summary command. Output fields are listed in the approximate order in which they appear. |

Table 97: show services service-sets summary Output Fields

| Field Name | Field Description |
|--------------------------------|--|
| Interface | Name of an adaptive services interface |
| Service type | Type of adaptive service, such as stateful firewall (SFW), Network Address Translation (NAT), intrusion detection service (IDS), Layer 2 Tunneling Protocol (L2TP), Compressed Real-Time Transport Protocol (CRTP), or IP Security (IPsec) |
| Service sets configured | Total number of service sets configured on the PIC that use internal service set IDs and do not consume external service sets, including CRTP and L2TP |
| Bytes used | Bytes used by a particular service or all services |
| Policy bytes used | Policy bytes used by a particular service or all services |
| CPU utilization | Percentage of the CPU resources being used |

Sample Output

show services service-sets summary

```

user@host> show services service-sets summary
Service sets
Interface  configured      Bytes used  Policy bytes used  CPU
utilization

```

| | | | | |
|----------|---|--------------------|------------------|-----|
| ms-4/0/0 | 1 | 14821556 (4.53 %) | 855124 (0.40 %) | N/A |
| ms-4/1/0 | 1 | 14691700 (4.49 %) | 855068 (0.40 %) | N/A |

show services service-sets summary interface

```
user@host> show services service-sets summary interface sp-1/3/0
Interface: sp-1/3/0
```

| Service type | Service sets
configured | Bytes used | CPU
utilization |
|--------------|----------------------------|------------------|--------------------|
| SFW/NAT/IDS | 1 | 54 (0.00 %) | N/A |
| L2TP | 1 | 58 (0.00 %) | N/A |
| C RTP | 1 | 58 (0.00 %) | N/A |
| System | 0 | 920831 (0.44 %) | N/A |
| Idle | 0 | 0 (0.00 %) | N/A |
| Total | 3 | 921001 (0.44 %) | N/A |

show services software

| | |
|---------------------------------|--|
| Syntax | show services software
<count> |
| Release Information | Command introduced in Junos OS Release 10.4.
<count> option added in Junos OS Release 11.2. |
| Description | Display information about software services. Information is displayed on both 6rd and DS-Lite services. |
| Options | count <i>interface-name</i> — (Optional) Display the current software counts for a service set for both DS-Lite and 6rd. |
| Required Privilege Level | view |
| List of Sample Output | show services software on page 2034
show services software count on page 2034 |
| Output Fields | Table 98 on page 2034 lists the output fields for the command-name command. Output fields are listed in the approximate order in which they appear. |

Table 98: show-services-software Output Fields

| Field Name | Field Description | Level of Output |
|--------------------|--|-----------------|
| Interface | Interface for which information is displayed. | All levels |
| Service Set | Service set containing the software rules for the interface. | All levels |
| Software | Name of the software concentrator. | All levels |
| Direction | Direction of the flow. | All levels |
| Flow count | Number of flows. | All levels |

Sample Output

show services software

```

user@host> show services software
Interface: sp-3/0/0, Service set: v6rd-dom1-dom3-service-set
Software          Direction    Flow count
10.10.10.2        ->          30.30.30.1    I           13

```

show services software count

```

user@host> show services software count
Interface  Service set    DS-Lite    6RD
sp-0/0/0   dslite-svc-set1  2          0

```

show services software flows

Syntax `show services software flows`
 (`<interface interface-name> <service-set service-set-name>|`
`count <interface interface-name> <service-set service-set-name>|`
`ds-lite <B4 b4-address> <AFTR aftr-address>|`
`v6rd <initiator initiator-ip-address><concentrator concentrator-ip-address>)`

Release Information Command introduced in Junos OS Release 10.2.

Description Display statistics information about the software flows.



NOTE: Starting with Junos OS Release 14.1R4, the IPv6 prefix length associated with a subscriber's basic broadband bridging device that is subject to a limited number of sessions (`dslite-ipv6-prefix-length` attribute) is taken into account while the session count is calculated and displayed in the output of the `show services software flows` command. Until Junos OS Release 14.1R3, only IPv4 flows were counted and IPv6 flows were not considered for the statistics about software flows

Options `interface interface-name`—(Optional) Display statistics information about the specified interface only.

`service-set service-set-name`—(Optional) Display statistics information about the specified service set only.

`count <interface interface-name> <service-set service-set-name>|`—(Optional) Display flow count information only, with optional filtering by interface and service set.

`ds-lite <B4 b4-address> <AFTR aftr-address>|`—(Optional) Display DS-Lite flow information, with optional filtering by B4 (software initiator) and AFTR (software concentrator).

`v6rd <initiator initiator-ip-address><concentrator concentrator-ip-address>|`—(Optional) Display v6rd flow information, with optional filtering by the software initiator and software concentrator.

Required Privilege Level view

List of Sample Output [show services software flows on page 2036](#)
[show services software flows count on page 2036](#)
[show services software flows ds-lite B4 on page 2037](#)
[show services software flows ds-lite AFTR on page 2037](#)
[services software flows ds-lite AFTR and B4 on page 2037](#)

Output Fields [Table 99 on page 2036](#) lists the output fields for the `show services software flows` command. Output fields are listed in the approximate order in which they appear.

Table 99: show services software flows Output Fields

| Field Name | Field Description |
|--------------------|--|
| Interface | Name of the interface. |
| Service set | Name of the service set. |
| Flow | Description of flow, including protocol input and output interface addresses. |
| State | Flow state. Value is: <ul style="list-style-type: none"> • Forward |
| Dir | Flow direction. Values are: <ul style="list-style-type: none"> • I—inbound • O—outbound |
| Frm count | Number of frames transferred. |
| NAT dest | NAT translation of the decapsulated address. |
| Software | For outbound flows, the address of the local software initiator (B4 for DS-Lite) is shown first, followed by the address of the software concentrator (AFTR for DS-Lite). For inbound flows, the address of the software concentrator is shown first, followed by the address of the software initiator. |

Sample Output

show services software flows

```

user@host> show services software flows
Interface: sp-0/0/0, Service set: dslite-svc-set1
Flow      State      Dir      Frm count
TCP       200.200.200.2:80 -> 33.33.33.1:1066 Forward 0      2005418
  NAT dest 33.33.33.1:1066 -> 20.20.1.2:1025
  Software 1001::1 -> 2001::2
TCP       20.20.1.2:1025 -> 200.200.200.2:80 Forward I      2007168
  NAT source 20.20.1.2:1025 -> 33.33.33.1:1066
  Software 2001::2 -> 1001::1
TCP       20.20.1.2:1025 -> 200.200.200.2:80 Forward I      2635998
  NAT source 20.20.1.2:1025 -> 33.33.33.1:1065
  Software 2001::3 -> 1001::1
DS-LITE   2001::2 -> 1001::1 Forward I      2008157
TCP       200.200.200.2:80 -> 33.33.33.1:1065 Forward O      2637909
  NAT dest 33.33.33.1:1065 -> 20.20.1.2:1025
  Software 1001::1 -> 2001::3
DS-LITE   2001::3 -> 1001::1 Forward I      2640499

```

show services software flows count

```

user@host> show services software flows count
Interface  Service set      Flow count
sp-0/0/0   dslite-svc-set1  6

```

show services software flows ds-lite B4

```

user@host> show services software flows ds-lite B4 2001::2
Interface: sp-0/0/0, Service set: dslite-svc-set1
Flow                                     State      Dir      Frm count
TCP      200.200.200.2:80  ->  33.33.33.1:1066 Forward  O      2884037
    NAT dest      33.33.33.1:1066  ->  20.20.1.2:1025
    Software      1001::1          ->  2001::2
TCP      20.20.1.2:1025  ->  200.200.200.2:80 Forward  I      2885884
    NAT source    20.20.1.2:1025  ->  33.33.33.1:1066
    Software      2001::2          ->  1001::1
DS-LITE   2001::2        ->  1001::1 Forward  I      2886821

```

show services software flows ds-lite AFTR

```

user@host> show services software flows ds-lite AFTR 1001::1
Interface: sp-0/0/0, Service set: dslite-svc-set1
Flow                                     State      Dir      Frm count
TCP      200.200.200.2:80  ->  33.33.33.1:1066 Forward  O      3359356
    NAT dest      33.33.33.1:1066  ->  20.20.1.2:1025
    Software      1001::1          ->  2001::2
TCP      20.20.1.2:1025  ->  200.200.200.2:80 Forward  I      3361235
    NAT source    20.20.1.2:1025  ->  33.33.33.1:1066
    Software      2001::2          ->  1001::1
TCP      20.20.1.2:1025  ->  200.200.200.2:80 Forward  I      4479810
    NAT source    20.20.1.2:1025  ->  33.33.33.1:1065
    Software      2001::3          ->  1001::1
DS-LITE   2001::2        ->  1001::1 Forward  I      3362168
TCP      200.200.200.2:80  ->  33.33.33.1:1065 Forward  O      4481520
    NAT dest      33.33.33.1:1065  ->  20.20.1.2:1025
    Software      1001::1          ->  2001::3
DS-LITE   2001::3        ->  1001::1 Forward  I      4484094

```

services software flows ds-lite AFTR and B4

```

user@host> show services software flows ds-lite AFTR 1001::1 B4 2001::2
Interface: sp-0/0/0, Service set: dslite-svc-set1
Flow                                     State      Dir      Frm count
TCP      200.200.200.2:80  ->  33.33.33.1:1066 Forward  O      3931026
    NAT dest      33.33.33.1:1066  ->  20.20.1.2:1025
    Software      1001::1          ->  2001::2
TCP      20.20.1.2:1025  ->  200.200.200.2:80 Forward  I      3932792
    NAT source    20.20.1.2:1025  ->  33.33.33.1:1066
    Software      2001::2          ->  1001::1
DS-LITE   2001::2        ->  1001::1 Forward  I      3933782

```

show services software statistics

| | |
|---------------------------------|--|
| Syntax | <pre>show services software statistics <ds-lite> <ds-lite> <interface interface-name> <v6rd></pre> |
| Release Information | Command introduced in Junos OS Release 10.4. |
| Description | Display information about software services. |
| Options | <p>ds-lite—(Optional) Display only DS-Lite.</p> <p>interface interface-name —(Optional) Name of the interface servicing the software. When you omit this option, data for all interfaces are shown.</p> <p>v6rd—(Optional) Display only 6rd statistics.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services software statistics on page 2041
show services software statistics ds-lite on page 2042 |
| Output Fields | Table 100 on page 2038 lists the output fields for the command-name command. Output fields are listed in the approximate order in which they appear. |

Table 100: command-name Output Fields

| Field Name | Field Description | Level of Output |
|------------------------------|--|-----------------------------|
| Service PIC Name | Name of service PIC for which statistics are shown. | statistics |
| Softwires Created | Number of softwires created. | statistics |
| Softwires Created for EIF/HP | Number of softwires created for endpoint-independent filtering (EIF) or hairpinning (HP). | statistics for ds-lite only |
| Softwires Deleted | Number of softwires deleted. | statistics |
| Softwires Flows Created | Number of flows created. | statistics |
| Softwires Flows Deleted | Number of flows deleted. | statistics |
| Slow Path Packets Processed | Number of packets processed as initial packets in a software session. These packets require a rule lookup and setting up of flows; this processing of an initial packet in a flow is called <i>the slow path</i> . | statistics |

Table 100: command-name Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|--|---------------------------------------|
| Slow Path Packets Processed for EIF/HP | Number of slow path EIF/HP packets processed. | statistics for ds-lite only |
| Fast Path Packets Processed | Number of packets processed that are not <i>slow path</i> . | statistics |
| Fast Path Encapsulated | Number of packets encapsulated in the fast path. | statistics |
| Softwire EIF Accept | Number of packets that matched an EIF entry that initiated the creation of a DS-Lite tunnel. The EIF entry was previously triggered by a DS-Lite packet. | statistics for ds-lite only |
| Rule Match Succeeded | Number of packets that matched a softwire rule. | statistics |
| Rule Match Failed | Number of packets that did not match any softwire rule. | statistics |
| IPv6 Packets Fragmented | Number of packets fragmented by the services PIC. | statistics for ds-lite only |
| IPv4 Client Fragments | Number of IPv4 fragments received from the client end over the softwire tunnel destined to the server. | statistics for ds-lite only |
| IPv4 Server First Fragments | Number of IPv4 first fragments received from the server destined to go over the softwire tunnel to the client. | statistics for ds-lite only |
| IPv4 Server More Fragments | Number of IPv4 other fragments (excluding first and last fragment) received from the server destined to go over the softwire tunnel to the client. | statistics for ds-lite only |
| IPv4 Server Last Fragments | Number of IPv4 last fragments received from the server destined to go over the softwire tunnel to the client. | statistics for ds-lite only |
| ICMPv4 Packets sent | Number of ICMPv4 packets sent to the softwire concentrator. | statistics |
| ICMPv4 Error Packets sent | Number of ICMPv4 error packets sent to the softwire concentrator. | statistics |
| ICMPv6 Packets sent | Number of ICMPv6 packets sent to the softwire concentrator. | statistics |
| Dropped ICMPv6 packets destined to AFTR | Number of ICMPv6 packets dropped instead of sending to the softwire concentrator. | statistics |
| Softwire Creation Failed | Number of softwire creation failures. | statistics for ds-lite and 6rd |

Table 100: command-name Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|--|------------------------------------|
| Softwire Creation Failed for EIF/HP | Number of softwire creation failures for EIF/HP. | statistics for ds-lite only |
| Flow Creation Failed | Number of flow creation failures. | statistics |
| Flow Creation Failed for EIF/HP | Number of flow creation failures for EIF/HP. | statistics for ds-lite only |
| Flow Creation Failed - Retry | Number of flow creations retried after failure. | statistics |
| Slow Path Failed | Number of failures detected in the slow path. | statistics |
| Slow Path Failed - Retry | Number of times processing of a packet was reprocessed in the slow path. | statistics |
| Packet not IPv4-in-IPv6 | Number of IPv4 packets not encapsulated in IPv6. | statistics for ds-lite only |
| IPv6 Fragmentation Error | Number of IPv6 packets with fragmentation errors. | statistics |
| Slow Path Failed-IPv6 Next Header Offset | Number of IPv6 header errors detected in slow path processing. | statistics for ds-lite only |
| Decapsulated Packet not IPv4 | Number of packets without IPv4 inner header. | statistics for ds-lite only |
| Decap Failed - IPv6 Next Header Offset | Decapsulation failure due to an unexpected inner header. | statistics for ds-lite only |
| Decap Failed - IPv4 L3 Integrity | Decapsulation failure due to incorrect Layer 3 data, such as not an IP packet, bad source or destination address, checksum error, or protocol error. | statistics for ds-lite only |
| Decap Failed - IPv4 L4 Integrity | Decapsulation failure due to incorrect Layer 4 data, such as errors in TCP, UDP, or TCP headers. | statistics for ds-lite only |
| No Softwire ID | Number of times a softwire ID was not found. | statistics |
| No Flow Extension | Number of times flow extensions were not found. | statistics |
| ICMPv4 Dropped Packets | Number of ICMPv4 packets dropped. | statistics |
| Packet not IPv6-in-IPv4 | Number of IPv6 packets not encapsulated in IPv4. | statistics for v6rd only |

Table 100: command-name Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|--|------------------------------------|
| Decapsulated
Packet not IPv6 | Number of packets without an IPv6 inner header. | statistics for v6rd only |
| Encapsulation
Failed - No packet
memory | Failed to encapsulate IPv6 packets in IPv4 due to low memory. | statistics for v6rd only |
| Flow limit exceeded | Flow not created because configured maximum flows per software is exceeded. | statistics |
| Session limit
exceeded | Flow not created because configured maximum DS-Lite software sessions per IPv6 prefix is exceeded. | statistics for ds-lite only |

Sample Output

show services software statistics

```
user@host> show services software statistics
DS-Lite Statistics:
```

```
Service PIC Name:                               :sp-0/0/0
```

Statistics

```
-----
```

```

Software Created                               :0
Software Created for EIF/HP                     :0
Software Deleted                               :0
Software Flows Created                         :0
Software Flows Deleted                         :0
Slow Path Packets Processed                     :0
Slow Path Packets Processed for EIF/HP          :0
Fast Path Packets Processed                     :0
Fast Path Packets Encapsulated                 :0
Software EIF Accept                            :0
Rule Match Succeeded                           :0
Rule Match Failed                             :0
IPv6 Packets Fragmented                       :0
IPv4 Client Fragments                          :0
IPv4 Server First Fragments                    :0
IPv4 Server More Fragments                     :0
IPv4 Server Last Fragments                     :0
ICMPv4 Packets sent                            :0
ICMPv4 Error Packets sent                      :0
ICMPv6 Packets sent                            :0
Dropped ICMPv6 packets destined to AFTR        :0
```

Transient Errors

```
-----
```

```

Flow Creation Failed - Retry                    :0
Flow Creation Failed - Retry for EIF/HP         :0
Slow Path Failed - Retry                        :0
```

Errors

| | |
|--|----|
| Softwire Creation Failed | :0 |
| Softwire Creation Failed for EIF/HP | :0 |
| Flow Creation Failed | :0 |
| Flow Creation Failed For EIF/HP | :0 |
| Slow Path Failed | :0 |
| Packet not IPv4-in-IPv6 | :0 |
| IPv6 Fragmentation Error | :0 |
| Softwire Creation Failed - IPv6 Next Header Offset | :0 |
| Decapsulated Packet not IPv4 | :0 |
| Decap Failed - IPv6 Next Header Offset | :0 |
| Decap Failed - IPv4 L3 Integrity | :0 |
| Decap Failed - IPv4 L4 Integrity | :0 |
| No Softwire ID | :0 |
| No Flow Extension | :0 |
| Flow Limit Exceeded | :0 |

6rd Statistics:

Service PIC Name :sp-0/0/0

Statistics

| | |
|--------------------------------|----|
| Softwires Created | :0 |
| Softwires Deleted | :0 |
| Softwires Flows Created | :0 |
| Softwires Flows Deleted | :0 |
| Slow Path Packets Processed | :0 |
| Fast Path Packets Processed | :0 |
| Fast Path Packets Encapsulated | :0 |
| Rule Match Failed | :0 |
| Rule Match Succeeded | :0 |

Transient Errors

| | |
|------------------------------|----|
| Flow Creation Failed - Retry | :0 |
| Slow Path Failed - Retry | :0 |

Errors

| | |
|--|----|
| Softwire Creation Failed | :0 |
| Flow Creation Failed | :0 |
| Slow Path Failed | :0 |
| Packet not IPv6-in-IPv4 | :0 |
| Slow Path Failed - IPv6 Next Header Offset | :0 |
| Decapsulated Packet not IPv6 | :0 |
| Encapsulation Failed - No packet memory | :0 |
| No Softwire ID | :0 |
| No Flow Extension | :0 |
| ICMPv4 Dropped Packets | :0 |

show services softwire statistics ds-lite

user@host> show services softwire statistics ds-lite

DS-Lite Statistics:

Service PIC Name: :sp-0/0/0

Statistics

| | |
|---|----|
| Softwires Created | :0 |
| Softwires Created for EIF/HP | :0 |
| Softwires Deleted | :0 |
| Softwires Flows Created | :0 |
| Softwires Flows Deleted | :0 |
| Slow Path Packets Processed | :0 |
| Slow Path Packets Processed for EIF/HP | :0 |
| Fast Path Packets Processed | :0 |
| Fast Path Packets Encapsulated | :0 |
| Software EIF Accept | :0 |
| Rule Match Succeeded | :0 |
| Rule Match Failed | :0 |
| IPv6 Packets Fragmented | :0 |
| IPv4 Client Fragments | :0 |
| IPv4 Server First Fragments | :0 |
| IPv4 Server More Fragments | :0 |
| IPv4 Server Last Fragments | :0 |
| ICMPv4 Packets sent | :0 |
| ICMPv4 Error Packets sent | :0 |
| ICMPv6 Packets sent | :0 |
| Dropped ICMPv6 packets destined to AFTR | :0 |

Transient Errors

| | |
|---|----|
| Flow Creation Failed - Retry | :0 |
| Flow Creation Failed - Retry for EIF/HP | :0 |
| Slow Path Failed - Retry | :0 |

Errors

| | |
|--|----|
| Software Creation Failed | :0 |
| Software Creation Failed for EIF/HP | :0 |
| Flow Creation Failed | :0 |
| Flow Creation Failed For EIF/HP | :0 |
| Slow Path Failed | :0 |
| Packet not IPv4-in-IPv6 | :0 |
| IPv6 Fragmentation Error | :0 |
| Software Creation Failed - IPv6 Next Header Offset | :0 |
| Decapsulated Packet not IPv4 | :0 |
| Decap Failed - IPv6 Next Header Offset | :0 |
| Decap Failed - IPv4 L3 Integrity | :0 |
| Decap Failed - IPv4 L4 Integrity | :0 |
| No Software ID | :0 |
| No Flow Extension | :0 |
| Flow Limit Exceeded | :0 |
| Session Limit Exceeded | :0 |

show services stateful-firewall conversations

Syntax show services stateful-firewall conversations
<brief | extensive | terse>
<application-protocol *protocol*>
<destination-port *destination-port*>
<destination-prefix *destination-prefix*>
<interface *interface-name*>
<limit *number*>
<pgcp>
<protocol *protocol*>
<service-set *service-set*>
<source-port *source-port*>
<source-prefix *source-prefix*>

Release Information Command introduced before Junos OS Release 7.4.
pgcp option introduced in Junos OS Release 8.4.

Description Display information about stateful firewall conversations.

Options **none**—Display standard information about all stateful firewall conversations.

brief | extensive | terse—(Optional) Display the specified level of output.

application-protocol *protocol*—(Optional) Display information about one of the following application protocols:

- **bootp**—Bootstrap protocol
- **dce-rpc**—Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—Domain Name System protocol
- **exec**—Exec
- **ftp**—File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol
- **iiop**—Internet Inter-ORB Protocol
- **login**—Login
- **netbios**—NetBIOS
- **netshow**—NetShow
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol
- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol

- **shell**—Shell
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

destination-port *destination-port*—(Optional) Display information for a particular destination port. The range of values is 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Display information for a particular destination prefix.

interface *interface-name*—(Optional) Display information about a particular interface. On M Series and T Series routers, the *interface-name* can be *sp-fpc/pic/port* or *rspnumber*.

limit *number*—(Optional) Maximum number of entries to display.

pgcp—(Optional) Display information about stateful firewall conversations for Packet Gateway Control Protocol (PGCP) flows.

protocol *protocol*—(Optional) Display information about one of the following IP types:

- **number**—Numeric protocol value from 0 to 255
- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Display information for the specific service set.

source-port *source-port*—(Optional) Display information for a particular source port. The range of values is 0 to 65535.

source-prefix *source-prefix*—(Optional) Display information for a particular source prefix.

Required Privilege Level view

List of Sample Output [show services stateful-firewall conversations on page 2047](#)
[show services stateful-firewall conversations destination-port on page 2047](#)

Output Fields Table 101 on page 2046 lists the output fields for the **show services stateful-firewall conversations** command. Output fields are listed in the approximate order in which they appear.

Table 101: show services stateful-firewall conversations Output Fields

| Field Name | Field Description |
|--------------------------|--|
| Interface | Name of an adaptive services interface. |
| Service set | Name of a service set. Individual empty service sets are not displayed, but if no service set has any flows, a flow table header is printed for each service set. |
| Conversation | Information about a group of related flows. <ul style="list-style-type: none"> • ALG Protocol—Application-level gateway protocol. • Number of initiators—Number of flows that initiated a session. • Number of responders—Number of flows that responded in a session. |
| Flow or Flow Prot | Protocol used for this flow. |
| Source | Source prefix of the flow, in the format <i>source-prefix-port</i> . |
| Destination | Destination prefix of the flow. |
| State | Status of the flow: <ul style="list-style-type: none"> • Drop—Drop all packets in the flow without response. • Forward—Forward the packet in the flow without looking at it. • Reject—Drop all packets in the flow with response. • Watch—Inspect packets in the flow. |
| Dir | Direction of the flow: input (I) or output (O). |
| Source NAT | Original and translated source IPv4 or IPv6 addresses are displayed if Network Address Translation (NAT) is configured on this particular flow or conversation. |
| Frm Count | Number of frames in the flow. |
| Destin NAT | Original and translated destination IPv4 or IPv6 addresses are displayed if NAT is configured on this particular flow or conversation. |

Table 101: show services stateful-firewall conversations Output Fields (*continued*)

| Field Name | Field Description |
|-----------------|---|
| Byte count | Number of bytes forwarded in the flow. |
| TCP established | Whether a TCP connection was established: Yes or No . |
| TCP window size | Negotiated TCP connection window size, in bytes. |
| TCP acknowledge | TCP acknowledgment sequence number. |
| TCP tickle | Whether TCP inquiry mode is on (enabled or disabled) and the time remaining to send the next inquiry, in seconds. |
| Master flow | Flow that initiated the conversation. |
| Timeout | Lifetime of the flow, in seconds. |

Sample Output

show services stateful-firewall conversations

```

user@host> show services stateful-firewall conversations
Interface: sp-1/3/0, Service set: green
Conversation: ALG Protocol: any, Number of initiators: 1,
Number of responders: 1

Flow
Prot      Source                Dest                State      Dir   Frm count
TCP       10.58.255.50:33005->    10.58.255.178:23   Forward    I     13
    Source NAT    10.58.255.50:33005->    10.59.16.100:4000
    Destin NAT    10.58.255.178:23 ->    0.0.0.0:4000
Byte count:          918
TCP established, TCP window size: 65535, TCP acknowledge: 2502627025
TCP tickle enabled, 0 seconds,
Master flow, Timeout: 30 seconds
TCP       10.58.255.178:23 ->    10.59.16.100:4000 Forward    0     8

```

show services stateful-firewall conversations destination-port

```

user@host> show services stateful-firewall conversations destination-port 21
Interface: sp-0/3/0, Service set: svc_set_trust

Interface: sp-0/3/0, Service set: svc_set_untrust
Conversation: ALG protocol: ftp
Number of initiators: 1, Number of responders: 1
Flow
TCP       10.50.10.2:2143 ->    10.50.20.2:21      Watch     0     0
TCP       10.50.20.2:21 ->    10.50.10.2:2143    Watch     I     0
TCP       10.50.20.2:21 ->    10.50.10.2:2143    Watch     I     0

```

show services stateful-firewall flow-analysis

| | |
|---------------------------------|--|
| Syntax | show services stateful-firewall flow-analysis
<interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 10.4R1. |
| Description | Display stateful firewall flow statistics. |
| Options | none —Display standard information about all stateful firewall flow statistics.

interface <i>interface-name</i> —(Optional) Display information about a particular interface. . |
| Required Privilege Level | view |
| List of Sample Output | show services stateful-firewall flow-analysis on page 2049
show services stateful-firewall flow-analysis interface sp-3/0/0 on page 2050 |
| Output Fields | Table 102 on page 2048 lists the output fields for the show services stateful-firewall flow-analysis command. Output fields are listed in the approximate order in which they appear. |

Table 102: show services stateful-firewall flow-analysis Output Fields

| Field Name | Field Description |
|------------------------------|---|
| Total Flows Active | Total active flows in the MS-PIC including TCP, UDP, ICMP and Softwires. |
| Total TCP Flows Active | Total active TCP flows in the MS-PIC. |
| Total UDP Flows Active | Total active UDP flows in the MS-PIC. |
| Total Other Flows Active | Total other active flows in the MS-PIC including ICMP and softwires. |
| Total Predicted Flows Active | Predicted flows are created only by the ALG traffic using the L3/L4 information available. |
| Created Flows per Second | Flow setup rate at the time of running the command. |
| Deleted Flows per Second | Flow deletion rate at the time of running the command. |
| Peak Total Flows Active | The highest number of active flows since the last PIC restart or since the last time flow statistics are flushed. |
| Peak Total TCP Flows Active | The highest number of active TCP flows since the last PIC restart or since the last time flow stats are flushed. |
| Peak Total UDP Flows Active | The highest number of active UDP flows since the last PIC restart or since the last time flow statistics are flushed. |

Table 102: show services stateful-firewall flow-analysis Output Fields (*continued*)

| Field Name | Field Description |
|---|---|
| Peak Total Other Flows Active | The highest number of other active flows since the last PIC restart or since the last time flow statistics are flushed. |
| Peak Created Flows per Second | The maximum flow setup rate observed since the last PIC restart or since the last time flow statistics are flushed. |
| Peak Deleted Flows per Second | The maximum flow deletion rate observed since the last PIC restart or from the last time flow statistics are flushed. |
| Average HTTP Flow Lifetime(ms) | Average HTTP Flow Lifetime in millisecond. |
| Packets received | The total number of packets received by the MS-PIC. |
| Packets transmitted | The total number of packets transmitted by the MS-PIC. |
| Slow path forward | The number of packets forwarded in the slow path (i.e. after the successful rule match and flow creation). |
| Slow path discard | The number of packets discarded before the flow creation. |
| Flow Rate Data: Number of Samples | The number of samples used to calculate the flow rate, since the last PIC restart or since the last time flow statistics are flushed. |
| Flow Rate Distribution(sec) Flow Operation :Creation Flow Operation :Deletion | Histogram of the samples used for flow rate calculation. |
| Flow Lifetime Distribution(sec): | Histogram of the samples used to calculate the flow life time in sec. |

Sample Output

show services stateful-firewall flow-analysis

```
user@host> show services stateful-firewall flow-analysis
```

```
Services PIC Name: sp-3/0/0
```

```
Flow Analysis Statistics:
```

```

Total Flows Active           :40
Total TCP Flows Active       :0
Total UDP Flows Active       :40
Total Other Flows Active     :0
Total Predicted Flows Active :0
Created Flows per Second     :0
Deleted Flows per Second     :0
Peak Total Flows Active      :40
Peak Total TCP Flows Active  :0
Peak Total UDP Flows Active  :40
Peak Total Other Flows Active :0
Peak Created Flows per Second :20
```

```

Peak Deleted Flows per Second      :20
Average HTTP Flow Lifetime(ms)     :0
Packets received                   :48682539117
Packets transmitted                 :48682502703
Slow path forward                   :6550
Slow path discard                   :0
Flow Rate Data:
Number of Samples: 19720
Flow Rate Distribution(sec)
Flow Operation :Creation
300000+          :0
250000 - 300000  :0
200000 - 250000  :0
160000 - 200000  :0
150000 - 160000  :0
50000 - 150000   :0
40000 - 50000    :0
30000 - 40000    :0
20000 - 30000    :0
10000 - 20000    :0
1000 - 10000     :0
0 - 1000         :19720
Flow Operation :Deletion
300000+          :0
250000 - 300000  :0
200000 - 250000  :0
160000 - 200000  :0
150000 - 160000  :0
50000 - 150000   :0
40000 - 50000    :0
30000 - 40000    :0
20000 - 30000    :0
10000 - 20000    :0
1000 - 10000     :0
0 - 1000         :19720
Flow Lifetime Distribution(sec):
      TCP      UDP      HTTP
240+      :0      0      0
120 - 240 :0      0
60 - 120  :0      0
30 - 60   :0      0
15 - 30   :0      6530
5 - 15    :0      0
1 - 5     :0      0
0 - 1     :0      6530

```

Sample Output

show services stateful-firewall flow-analysis interface sp-3/0/0

```

user@host> show services stateful-firewall flow-analysis interface sp-3/0/0
Services PIC Name: sp-3/0/0
Flow Analysis Statistics:
Total Flows Active          :40
Total TCP Flows Active      :0
Total UDP Flows Active      :40
Total Other Flows Active    :0
Total Predicted Flows Active :0
Created Flows per Second    :0
Deleted Flows per Second    :0
Peak Total Flows Active     :40

```

```

Peak Total TCP Flows Active      :0
Peak Total UDP Flows Active     :40
Peak Total Other Flows Active   :0
Peak Created Flows per Second   :20
Peak Deleted Flows per Second   :20
Average HTTP Flow Lifetime(ms) :0
Packets received                :54696856768
Packets transmitted             :54696815873
Slow path forward               :7350
Slow path discard               :0
Flow Rate Data:
Number of Samples: 22139
Flow Rate Distribution(sec)
Flow Operation :Creation
300000+        :0
250000 - 300000 :0
200000 - 250000 :0
160000 - 200000 :0
150000 - 160000 :0
50000 - 150000  :0
40000 - 50000   :0
30000 - 40000   :0
20000 - 30000   :0
10000 - 20000   :0
1000 - 10000    :0
0 - 1000        :22139
Flow Operation :Deletion
300000+        :0
250000 - 300000 :0
200000 - 250000 :0
160000 - 200000 :0
150000 - 160000 :0
50000 - 150000  :0
40000 - 50000   :0
30000 - 40000   :0
20000 - 30000   :0
10000 - 20000   :0
1000 - 10000    :0
0 - 1000        :22139
Flow Lifetime Distribution(sec):
      TCP      UDP      HTTP
240+      :0      0      0
120 - 240 :0      0      0
60 - 120  :0      0      0
30 - 60   :0      0      0
15 - 30   :0      7330   0
5 - 15    :0      0      0
1 - 5     :0      0      0
0 - 1     :0      7330   0

```

show services stateful-firewall flows

Syntax `show services stateful-firewall flows`
 `<brief | extensive | summary | terse>`
 `<application-protocol protocol>`
 `<count>`
 `<destination-port destination-port>`
 `<destination-prefix destination-prefix>`
 `<interface interface-name>`
 `<limit number>`
 `<protocol protocol>`
 `<service-set service-set>`
 `<source-port source-port>`
 `<source-prefix source-prefix>`

Release Information Command introduced before Junos OS Release 7.4.
 pgcp option introduced in Junos OS Release 8.4.
 application-protocol option introduced in Junos OS Release 10.4.

Description Display stateful firewall flow table entries. When the interface is used for software processing, the type of software concentrator (**DS-LITE** or **6rd**) is shown, and frame counts are provided.

Options **none**—Display standard information about all stateful firewall flows.

brief | extensive | summary | terse—(Optional) Display the specified level of output.

application-protocol *application-protocol*—(Optional) Display information about one of the following application-level gateway (ALG) protocol types:

- **bootp**—Bootstrap protocol
- **dce-rpc**—Distributed Computing Environment (DCE) remote procedure call (RPC) protocol



NOTE: Use this option to select Microsoft Remote Procedure Call (MSRPC).

- **dce-rpc-portmap**—Distributed Computing Environment (DCE) remote procedure call (RPC) portmap protocol
- **dns**—Domain Name Service protocol
- **exec**—Remote execution protocol
- **ftp**—File Transfer Protocol
- **h323**—H.323 protocol
- **icmp**—Internet Control Message Protocol
- **iioip**—Internet Inter-ORB Protocol

- **ip**—Internet protocol
- **netbios**—NetBIOS protocol
- **netshow**—Netshow protocol
- **pptp**—Point-to-Point Tunneling Protocol
- **realaudio**—RealAudio protocol
- **rpc**—Remote Procedure Call protocol



NOTE: Use this option to select Sun Microsystems Remote Procedure Call protocol (SunRPC).

- **rpc-portmap**—Remote Procedure Call portmap protocol
- **rtsp**—Real-Time Streaming Protocol
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **talk**—Talk protocol
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

count—(Optional) Display a count of the matching entries.

destination-port *destination-port*—(Optional) Display information for a particular destination port. The range of values is from 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Display information for a particular destination prefix.

interface *interface-name*—(Optional) Display information about a particular interface. On M Series and T Series routers, *interface-name* can be *ms-fpc/pic/port* or *rspnumber*.

limit *number*—(Optional) Maximum number of entries to display.

protocol *protocol*—(Optional) Display information about one of the following IP types:

- ***number***—Numeric protocol value from 0 to 255
- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol

- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Display information for a particular service set.

source-port *source-port*—(Optional) Display information for a particular source port. The range of values is from 0 to 65535.

source-prefix *source-prefix*—(Optional) Display information for a particular source prefix.

Required Privilege Level view

Related Documentation • [clear services stateful-firewall flows on page 1854](#)

List of Sample Output [show services stateful-firewall flows on page 2055](#)
[show services stateful-firewall flows \(For Software Flows\) on page 2055](#)
[show services stateful-firewall flows brief on page 2056](#)
[show services stateful-firewall flows extensive on page 2056](#)
[show services stateful-firewall flows count on page 2056](#)
[show services stateful-firewall flows destination port on page 2056](#)
[show services stateful-firewall flows source port on page 2056](#)
[show services stateful-firewall flows \(Twice NAT\) on page 2056](#)

Output Fields [Table 103 on page 2054](#) lists the output fields for the **show services stateful-firewall flows** command. Output fields are listed in the approximate order in which they appear.

Table 103: show services stateful-firewall flows Output Fields

| Field Name | Field Description |
|--------------------------|---|
| Interface | Name of the interface. |
| Service set | Name of a service set. Individual empty service sets are not displayed. If no service set has any flows, a flow table header is displayed for each service set. |
| Flow Count | Number of flows in a session. |
| Flow or Flow Prot | Protocol used for this flow. |
| Source | Source prefix of the flow in the format <i>source-prefix:port</i> . For ICMP flows, port information is not displayed. |

Table 103: show services stateful-firewall flows Output Fields (*continued*)

| Field Name | Field Description |
|------------------|--|
| Dest | Destination prefix of the flow. For ICMP flows, port information is not displayed. |
| State | Status of the flow: <ul style="list-style-type: none"> • Drop—Drop all packets in the flow without response. • Forward—Forward the packet in the flow without looking at it. • Reject—Drop all packets in the flow with response. • Watch—Inspect packets in the flow. |
| Dir | Direction of the flow: input (I) or output (O). |
| Frm count | Number of frames in the flow. |

Sample Output

show services stateful-firewall flows

```
user@host> show services stateful-firewall flows
Interface: ms-1/3/0, Service set: green
```

```
Flow
Prot      Source          Dest              State    Dir    Frm count
TCP       10.58.255.178:23 -> 10.59.16.100:4000 Forward  O
TCP       10.58.255.50:33005-> 10.58.255.178:23 Forward  I      1
Source NAT 10.58.255.50:33005-> 10.59.16.100:4000
Destin NAT 10.58.255.178:23 -> 0.0.0.0:4000
```

show services stateful-firewall flows (For Software Flows)

When a service set includes software processing, the following output format is used for the software flows:

```
user@host> show services stateful-firewall flows
Interface: sp-0/1/0, Service set: dslite-svc-set2
Flow
TCP       200.200.200.2:80 -> 44.44.44.1:1025 Forward  O      219942
NAT dest  44.44.44.1:1025 -> 20.20.1.4:1025
Software  2001::2 -> 1001::1
TCP       20.20.1.2:1025 -> 200.200.200.2:80 Forward  I      110244
NAT source 20.20.1.2:1025 -> 44.44.44.1:1024
Software  2001::2 -> 1001::1
TCP       200.200.200.2:80 -> 44.44.44.1:1024 Forward  O      219140
NAT dest  44.44.44.1:1024 -> 20.20.1.2:1025
Software  2001::2 -> 1001::1
DS-LITE   2001::2 -> 1001::1 Forward  I      988729
TCP       200.200.200.2:80 -> 44.44.44.1:1026 Forward  O      218906
NAT dest  44.44.44.1:1026 -> 20.20.1.3:1025
Software  2001::2 -> 1001::1
TCP       20.20.1.3:1025 -> 200.200.200.2:80 Forward  I      110303
NAT source 20.20.1.3:1025 -> 44.44.44.1:1026
Software  2001::2 -> 1001::1
TCP       20.20.1.4:1025 -> 200.200.200.2:80 Forward  I      110944
```

```

NAT source      20.20.1.4:1025  ->    44.44.44.1:1025
Software        2001::2         ->    1001::1

```

show services stateful-firewall flows brief

The output for the **show services stateful-firewall flows brief** command is identical to that for the **show services stateful-firewall flows** command. For sample output, see [show services stateful-firewall flows](#).

show services stateful-firewall flows extensive

```

user@host> show services stateful-firewall flows extensive
Interface: ms-0/3/0, Service set: ss_nat
Flow
count
TCP      16.1.0.1:2330  ->    16.49.0.1:21      Forward  I
8
  NAT source      16.1.0.1:2330  ->    16.41.0.1:2330
  NAT dest        16.49.0.1:21  ->    16.99.0.1:21
  Byte count: 455, TCP established, TCP window size: 57344
  TCP acknowledge: 3251737524, TCP tickle enabled, tcp_tickle: 0
  Flow role: Master, Timeout: 720
TCP      16.99.0.1:21   ->    16.41.0.1:2330     Forward  0
5
  NAT source      16.99.0.1:21   ->    16.49.0.1:21
  NAT dest        16.41.0.1:2330 ->    16.1.0.1:2330
  Byte count: 480, TCP established, TCP window size: 57344
  TCP acknowledge: 463128048, TCP tickle enabled, tcp_tickle: 0
  Flow role: Responder, Timeout: 720

```

show services stateful-firewall flows count

```

user@host> show services stateful-firewall flows count
Interface      Service set      Flow Count
ms-1/3/0       green            2

```

show services stateful-firewall flows destination port

```

user@router> show services stateful-firewall flows destination-port 21
Interface: ms-0/3/0, Service set: svc_set_trust
Flow
Interface: ms-0/3/0, Service set: svc_set_untrust
Flow
TCP      10.50.10.2:2143  ->    10.50.20.2:21      Watch   Dir    Frm count
                                State   Dir    Frm count
                                0      0      0

```

show services stateful-firewall flows source port

```

user@router> show services stateful-firewall flows source-port 2143
Interface: ms-0/3/0, Service set: svc_set_trust
Flow
Interface: ms-0/3/0, Service set: svc_set_untrust
Flow
TCP      10.50.10.2:2143  ->    10.50.20.2:21      Watch   Dir    Frm count
                                State   Dir    Frm count
                                0      0      0

```

show services stateful-firewall flows (Twice NAT)

```

user@router> show services stateful-firewall flows

```

| Flow | | | State | Dir | Frm count |
|------|-------------------|----------------------|-------------------|-----|-----------|
| UDP | 40.0.0.8:23439 | -> 80.0.0.1:16485 | Watch | I | 20 |
| | NAT source | 40.0.0.8:23439 -> | 172.16.1.10:1028 | | |
| | NAT dest | 80.0.0.1:16485 -> | 192.16.1.10:22415 | | |
| UDP | 192.16.1.10:22415 | -> 172.16.1.10:1028 | Watch | O | 20 |
| | NAT source | 192.16.1.10:22415 -> | 80.0.0.1:16485 | | |
| | NAT dest | 172.16.1.10:1028 -> | 40.0.0.8:23439 | | |

show services stateful-firewall sip-call

Syntax show services stateful-firewall sip-call
 <brief | extensive | terse>
 <application-protocol *protocol*>
 <destination-port *destination-port*>
 <destination-prefix *destination-prefix*>
 <interface *interface-name*>
 <limit *number*>
 <protocol *protocol*>
 <service-set *service-set*>
 <source-port *source-port*>
 <source-prefix *source-prefix*>

Release Information Command introduced in Junos OS Release 7.4.

Description Display stateful firewall Session Initiation Protocol (SIP) call information.

Options **count**—(Optional) Display a count of the matching entries.

brief—(Optional) Display brief SIP call information.

extensive—(Optional) Display detailed SIP call information.

terse—(Optional) Display terse SIP call information.

application-protocol—(Optional) Display information about one of the following application protocols:

- **bootp**—(SIP only) Bootstrap protocol
- **dce-rpc**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—(SIP only) Domain Name System protocol
- **exec**—(SIP only) Exec
- **ftp**—(SIP only) File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol
- **iiop**—Internet Inter-ORB Protocol
- **login**—Login
- **netbios**—NetBIOS
- **netshow**—NetShow
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol

- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **shell**—Shell
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

destination-port *destination-port*—(Optional) Display information for a particular destination port. The range of values is from 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Display information for a particular destination prefix.

interface *interface-name*—(Optional) Display information about a particular adaptive services interface. On M Series and T Series routers, *interface-name* can be *sp-fpc/pic/port* or *rspnumber*.

limit *number*—(Optional) Maximum number of entries to display.

protocol—(Optional) Display information about one of the following IP types:

- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol
- **ipv6**—IPv6 within IP
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Display information for a particular service set.

source-port *source-port*—(Optional) Display information for a particular source port. The range of values is from 0 to 65535.

source-prefix *source-prefix*—(Optional) Display information for a particular source prefix.

Required Privilege Level view

Related Documentation • [clear services stateful-firewall sip-call on page 1856](#)

List of Sample Output [show services stateful-firewall sip-call extensive on page 2061](#)

Output Fields [Table 104 on page 2060](#) lists the output fields for the **show services stateful-firewall sip-call** command. Output fields are listed in the approximate order in which they appear.

Table 104: show services stateful-firewall sip-call Output Fields

| Field Name | Field Description |
|----------------------------------|---|
| Interface | Name of an adaptive services interface. |
| Service set | Name of a service set. |
| From | Initiator address. |
| To | Responder address. |
| Call ID | SIP call identification string. |
| Number of initiator flows | Number of control , contact , or media initiator flows. |
| Number of responder flows | Number of control , contact , or media responder flows. |
| protocol | Protocol used for this flow. |
| source-prefix | Source prefix of the flow in the format source-prefix : port . |
| destination-prefix | Destination prefix of the flow. |
| state | Status of the flow: <ul style="list-style-type: none"> • Drop—Drop all packets in the flow without a response. • Forward—Forward the packet in the flow without examining it. • Reject—Drop all packets in the flow with a response. • Unknown—Unknown status. • Watch—Inspect packets in the flow. |
| direction | Direction of the flow: input (I), output (O), or unknown (U). |

Table 104: show services stateful-firewall sip-call Output Fields (*continued*)

| Field Name | Field Description |
|--------------------|--|
| <i>frame-count</i> | Number of frames in the flow. |
| Byte count | Number of bytes forwarded in the flow. |
| Flow role | Role of the flow that is under evaluation: Initiator , Master , Responder , or Unknown . |
| Timeout | Lifetime of the flow, in seconds. |

Sample Output

show services stateful-firewall sip-call extensive

```

user@host> show services stateful-firewall sip-call extensive
Interface: sp-0/3/0, Service set: test_sip_777

From : 6507771234@10.200.100.1:0;000ff73ac89900021bb231dc-3ef68435
To : 4085551234@10.200.100.1:0;0011bb65c2a3000777bd0fc-5748b749
Call ID : 000ff73a-c8990004-0741adac-3e027c7e@10.20.70.2
Number of control initiator flows: : 1, Number of control responder flows:
: 1
UDP      10.20.70.2:50354 -> 10.200.100.1:5060 Watch    I
2
  Byte count: 1112
  Flow role: Master, Timeout: 30
UDP      10.200.100.1:5060 -> 10.20.170.111:50354 Watch    0
0
  Byte count: 0
  Flow role: Responder, Timeout: 30
UDP      0.0.0.0:0 -> 10.20.170.111:5060 Watch    0
7
  Byte count: 2749
  Flow role: Responder, Timeout: 30
Number of contact initiator flows: 1, Number of contact responder flows: 1
UDP      0.0.0.0:0 -> 10.20.140.11:5060 Watch    I
1
  Byte count: 409
  Flow role: Master, Timeout: 30
UDP      10.20.140.11:31864 -> 10.20.170.111:18808 Forward  0
622
  Byte count: 124400
  Flow role: Master, Timeout: 30
UDP      0.0.0.0:0 -> 10.20.170.111:18809 Forward  0
0
  Byte count: 0
  Flow role: Initiator, Timeout: 30
Number of media initiator flows: 4, Number of media responder flows: 0
UDP      10.20.70.2:18808 -> 10.20.140.11:31864 Forward  I
628
  Byte count: 125600
  Flow role: Initiator, Timeout: 30
UDP      0.0.0.0:0 -> 10.20.140.11:31865 Forward  I
0

```

```
Byte count: 0
Flow role: Initiator, Timeout: 30
0          0.0.0.0:0    ->      0.0.0.0:0    Unknown  U
0
Byte count: 0
Flow role: Unknown, Timeout: 0
0          0.0.0.0:0    ->      0.0.0.0:0    Unknown  U
Interface: sp-0/3/0, Service set: test_sip_888
```


show services stateful-firewall sip-register

Syntax show services stateful-firewall sip-register
 <brief | extensive | terse>
 <application-protocol *protocol*>
 <destination-port *destination-port*>
 <destination-prefix *destination-prefix*>
 <interface *interface-name*>
 <limit *number*>
 <protocol *protocol*>
 <service-set *service-set*>
 <source-port *source-port*>
 <source-prefix *source-prefix*>

Release Information Command introduced in Junos OS Release 7.4.

Description Display stateful firewall Session Initiation Protocol (SIP) register information.

Options **count**—(Optional) Display a count of the matching entries.

brief—(Optional) Display brief SIP register information.

extensive—(Optional) Display detailed SIP register information.

terse—(Optional) Display terse SIP register information.

application-protocol—(Optional) Display information about one of the following application protocols:

- **bootp**—(SIP only) Bootstrap protocol
- **dce-rpc**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—(SIP only) Domain Name System protocol
- **exec**—(SIP only) Exec
- **ftp**—(SIP only) File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol
- **iiop**—Internet Inter-ORB Protocol
- **login**—Login
- **netbios**—NetBIOS
- **netshow**—NetShow
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol

- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **shell**—Shell
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

destination-port *destination-port*—(Optional) Display information for a particular destination port.

destination-prefix *destination-prefix*—(Optional) Display information for a particular destination prefix. The range of values is from 0 to 65535.

interface *interface-name*—(Optional) Display information about a particular interface. On M Series and T Series routers, the *interface-name* can be *sp-fpc/pic/port* or *rspnumber*.

limit *number*—(Optional) Maximum number of entries to display.

protocol—(Optional) Display information about one of the following IP types:

- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol
- **ipv6**—IPv6 within IP
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Display information for a particular service set.

source-port *source-port*—(Optional) Display information for a particular source port. The range of values is from 0 to 65535.

source-prefix *source-prefix*—(Optional) Display information for a particular source prefix.

Required Privilege Level view

Related Documentation • [clear services stateful-firewall sip-register on page 1859](#)

List of Sample Output [show services stateful-firewall sip-register extensive on page 2065](#)

Output Fields [Table 105 on page 2065](#) lists the output fields for the **show services stateful-firewall sip-register** command. Output fields are listed in the approximate order in which they appear.

Table 105: show services stateful-firewall sip-register Output Fields

| Field Name | Field Description |
|---------------------------|---|
| Interface | Name of an adaptive services interface. |
| Service set | Name of a service set. |
| SIP Register | Register information header. |
| Protocol | Protocol used for this flow. |
| Registered IP | Register IP address. |
| Port | Register port number. |
| Expiration timeout | Configured lifetime, in seconds. |
| Timeout remaining | Lifetime remaining, in seconds. |
| From | Initiator address. |
| To | Responder address. |
| Call ID | SIP call identification string. |

Sample Output

[show services stateful-firewall sip-register extensive](#)

```
user@host> show services stateful-firewall sip-register extensive
Interface: sp-0/3/0, Service set: test_sip_777
```

```
SIP Register: Protocol: UDP, Registered IP: 10.20.170.111, Port: 5060, Acked
Expiration timeout: 36000, Timeout remaining: 35544
From: : 6507771234@10.200.100.1:0;
```

To: : 6507771234@10.200.100.1:0;
Call ID: : 000ff73a-c8990002-23b1d942-2ba1f91f@10.20.70.2

Interface: sp-0/3/0, Service set: test_sip_888

SIP Register: Protocol: UDP, Registered IP: 10.20.170.112, Port: 5060, Acked
Expiration timeout: 36000, Timeout remaining: 35549
From: : 8881234@10.200.100.1:0;
To: : 8881234@10.200.100.1:0;
Call ID: : 00112096-81fc0002-23b38905-7cb41f62@10.20.71.2

show services stateful-firewall statistics

| | |
|---------------------------------|--|
| Syntax | show services stateful-firewall statistics
<application-protocol <i>protocol</i> >
<brief detail extensive summary>
<interface <i>interface-name</i> >
<service-set <i>service-set</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display stateful firewall statistics. |
| Options | <p>none—Display standard information about all stateful firewall statistics.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display information about a particular interface.
On M Series and T Series routers, the <i>interface-name</i> can be <i>ms-fpc/pic/port</i> or <i>rspnumber</i>.</p> <p>service-set <i>service-set</i>—(Optional) Display information about a particular service set.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear services stateful-firewall statistics on page 1862 |
| List of Sample Output | show services stateful-firewall statistics extensive on page 2074 |
| Output Fields | Table 106 on page 2067 lists the output fields for the show services stateful-firewall statistics command. Output fields are listed in the approximate order in which they appear. |

Table 106: show services stateful-firewall statistics Output Fields

| Field Name | Field Description |
|--|--|
| Interface | Name of an adaptive services interface. |
| Service set | Name of a service set. |
| New flows | Rule match counters for new flows: <ul style="list-style-type: none"> Rule Accepts—New flows accepted. Rule Discards—New flows discarded. Rule Rejects—New flows rejected. |
| Existing flow types packet counters | Rule match counters for existing flows: <ul style="list-style-type: none"> Accepts—Match existing forward or watch flow. Drop—Match existing discard flow. Rejects—Match existing reject flow. |

Table 106: show services stateful-firewall statistics Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------------|--|
| Hairpinning Counters | <p>Hairpinning counters:</p> <ul style="list-style-type: none"> • Slow Path Hairpinned Packets—Slow path packets that were hairpinned back to the internal network. • Fast Path Hairpinned Packets—Fast path packets that were hairpinned back to the internal network. |
| Drops | <p>Drop counters:</p> <ul style="list-style-type: none"> • IP option—Packets dropped in IP options processing. • TCP SYN defense—Packets dropped by SYN defender. • NAT ports exhausted—Hide mode. The router has no available Network Address Translation (NAT) ports for a given address or pool. • Sessions dropped due to subscriber flow limit—Sessions dropped because the subscriber's flow limit was exceeded. |
| Errors | <p>Total errors, categorized by protocol:</p> <ul style="list-style-type: none"> • IP—Total IP version 4 errors. • TCP—Total Transmission Control Protocol (TCP) errors. • UDP—Total User Datagram Protocol (UDP) errors. • ICMP—Total Internet Control Message Protocol (ICMP) errors. • Non-IP packets—Total non-IPv4 errors. • ALG—Total application-level gateway (ALG) errors |

Table 106: show services stateful-firewall statistics Output Fields (*continued*)

| Field Name | Field Description |
|------------|---|
| IP Errors | <p>IPv4 errors:</p> <ul style="list-style-type: none"> • IP packet length inconsistencies—IP packet length does not match the Layer 2 reported length. • Minimum IP header length check failures—Minimum IP header length is 20 bytes. The received packet contains less than 20 bytes. • Reassembled packet exceeds maximum IP length—After fragment reassembly, the reassembled IP packet length exceeds 65,535. • Illegal source address 0—Source address is not a valid address. Invalid addresses are, loopback, broadcast, multicast, and reserved addresses. Source address 0, however, is allowed to support BOOTP and the destination address 0xffffffff. • Illegal destination address 0—Destination address is not a valid address. The address is reserved. • TTL zero errors—Received packet had a time-to-live (TTL) value of 0. • Illegal IP protocol number (0 or 255)—IP protocol is 0 or 255. • Land attack—IP source address is the same as the destination address. • Non-IPv4 packets—Packet was not IPv4. (Only IPv4 is supported.) • Bad checksum—Packet had an invalid IP checksum. • Illegal IP fragment length—Illegal fragment length. All fragments (other than the last fragment) must have a length that is a multiple of 8 bytes. • IP fragment overlap—Fragments have overlapping fragment offsets. • IP fragment reassembly timeout—Some of the fragments for an IP packet were not received in time, and the reassembly handler dropped partial fragments. • IP fragment limit exceeded: 0—Fragments that exceeded the limit. • Unknown: 0—Unknown fragments. |

Table 106: show services stateful-firewall statistics Output Fields (*continued*)

| Field Name | Field Description |
|------------|-------------------|
| TCP Errors | |

Table 106: show services stateful-firewall statistics Output Fields (*continued*)

| Field Name | Field Description |
|------------|---|
| | TCP protocol errors: |
| | <ul style="list-style-type: none"> • TCP header length inconsistencies—Minimum TCP header length is 20 bytes, and the IP packet received does not contain at least 20 bytes. • Source or destination port number is zero—TCP source or destination port is zero. • Illegal sequence number and flags combinations — Dropped because of TCP errors, such as an illegal sequence number, which causes an illogical combination of flags to be set. • SYN attack (multiple SYN messages seen for the same flow)—Multiple SYN packets received for the same flow are treated as a SYN attack. The packets might be retransmitted SYN packets and therefore valid, but a large number is cause for concern. • First packet not a SYN message—First packets for a connection are not SYN packets. These packets might originate from previous connections or from someone performing an ACK/FIN scan. • TCP port scan (TCP handshake, RST seen from server for SYN)—In the case of a SYN defender, if an RST (reset) packet is received instead of a SYN/ACK message, someone is probably trying to scan the server. This behavior can result in false alarms if the RST packet is not combined with an intrusion detection service (IDS). • Bad SYN cookie response—SYN cookie generates a SYN/ACK message for all incoming SYN packets. If the ACK received for the SYN/ACK message does not match, this counter is incremented. • TCP reconstructor sequence number error—This counter is incremented in the following cases:
The TCP seqno is 0 and all the TCP flags are also 0.
The TCP seqno is 0 and FIN/PSH/URG TCP flags are set. • TCP reconstructor retransmissions—This counter is incremented for the retransmitted packets during connection 3-way handshake. • TCP partially opened connection timeout (SYN)—This counter is incremented when the SYN Defender is enabled and the 3-way handshake is not completed within the SYN DEFENDER TIMEOUT. The connection will be closed and resources will be released by sending RST to the responder. • TCP partially opened connection timeout (SYN-ACK)—This counter is incremented when the SYN Defender is enabled and the 3-way handshake is not completed within the SYN DEFENDER TIMEOUT. The connection will be closed and resources will be released by sending RST to the responder. • TCP partially closed connection reuse—Not supported. • TCP 3-way error - client sent SYN+ACK—A SYN/ACK should be sent by the server on receiving a SYN. This counter is incremented when the first message received from the initiator is SYN+ACK. • TCP 3-way error - server sent ACK—ACK should be sent by the client on receiving a SYN/ACK from the server. This counter is incremented when the ACK is received from the Server instead of from the Client. • TCP 3-way error - SYN seq number retransmission mismatch—This counter is incremented when the SYN is received again with a different sequence number from the first SYN sequence number. • TCP 3-way error - RST seq number mismatch—A reset could be received from either side. The server could send a RST on receiving a SYN or the client could send a RST on receiving SYN/ACK. This counter is incremented when the |

Table 106: show services stateful-firewall statistics Output Fields (*continued*)

| Field Name | Field Description |
|-------------|--|
| | <p>RST is received either from the client or server with a non-matching sequence number.</p> <ul style="list-style-type: none"> • TCP 3-way error - FIN received—This counter is incremented when the FIN is received during the 3-way handshake. • TCP 3-way error - invalid flags (PSH, URG, ECE, CWR)—This counter is incremented when any of the PSH, URG, ECE, or CWR flags were received during the 3-way handshake. • TCP 3-way error - SYN recvd but no client flows—This counter is incremented when SYN is received but not from the connection initiator. The counter is not incremented in the case of simultaneous open, when the SYN is received in both the directions. • TCP 3-way error - first packet SYN+ACK—The first packet received was SYN+ACK instead of SYN. • TCP 3-way error - first packet FIN+ACK—The first packet received was FIN+ACK instead of SYN. • TCP 3-way error - first packet FIN—The first packet received was FIN instead of SYN. • TCP 3-way error - first packet RST—The first packet received was RST instead of SYN. • TCP 3-way error - first packet ACK—The first packet received was ACK instead of SYN. • TCP 3-way error - first packet invalid flags (PSH, URG, ECE, CWR)—The first packet received had invalid flags. • TCP Close error - no final ACK—This counter is incremented when ACK is not received after the FINs are received from both directions. • TCP Resumed Flow—Plain ACKs create flows if rule match permits, and these are classified as TCP Resumed Flows. This counter is incremented in the case of a TCP Resumed Flow. |
| UDP Errors | <p>UDP protocol errors:</p> <ul style="list-style-type: none"> • IP data length less than minimum UDP header length (8 bytes)—Minimum UDP header length is 8 bytes. The received IP packets contain less than 8 bytes. • Source or destination port is zero—UDP source or destination port is 0. • UDP port scan (ICMP error seen for UDP flow)—ICMP error is received for a UDP flow. This could be a genuine UDP flow, but it is counted as an error. |
| ICMP Errors | <p>ICMP protocol errors:</p> <ul style="list-style-type: none"> • IP data length less than minimum ICMP header length (8 bytes)—ICMP header length is 8 bytes. This counter is incremented when received IP packets contain less than 8 bytes. • ICMP error length inconsistencies—Minimum length of an ICMP error packet is 48 bytes, and the maximum length is 576 bytes. This counter is incremented when the received ICMP error falls outside this range. • Duplicate ping sequence number—Received ping packet has a duplicate sequence number. • Mismatched ping sequence number—Received ping packet has a mismatched sequence number. • No matching flow—No matching existing flow was found for the ICMP error. |

Table 106: show services stateful-firewall statistics Output Fields (*continued*)

| Field Name | Field Description |
|-------------------|---|
| ALG errors | <p>Accumulation of all the application-level gateway protocol (ALG) drops counted separately in the ALG context:</p> <ul style="list-style-type: none"> • BOOTP—Bootstrap protocol errors • DCE-RPC—Distributed Computing Environment-Remote Procedure Call protocols errors • DCE-RPC portmap—Distributed Computing Environment-Remote Procedure Call protocols portmap service errors • DNS—Domain Name System protocol errors • Exec—Exec errors • FTP—File Transfer Protocol errors • H323—H.323 standards errors • ICMP—Internet Control Message Protocol errors • IIOB—Internet Inter-ORB Protocol errors • Login—Login errors • NetBIOS—NetBIOS errors • Netshow—NetShow errors • Real Audio—RealAudio errors • RPC—Remote Procedure Call protocol errors • RPC portmap—Remote Procedure Call protocol portmap service errors • RTSP—Real-Time Streaming Protocol errors • Shell—Shell errors • SIP—Session Initiation Protocol errors • SNMP—Simple Network Management Protocol errors • SQLNet—SQLNet errors • TFTP—Trivial File Transfer Protocol errors • Traceroute—Traceroute errors |
| Drop Flows | <ul style="list-style-type: none"> • Maximum Ingress Drop flows allowed—Maximum number of ingress flow drops allowed. • Maximum Egress Drop flows allowed—Maximum number of egress flow drops allowed. • Current Ingress Drop flows—Current number of ingress flow drops. • Current Egress Drop flows—Current number of egress flow drops. • Ingress Drop Flow limit drops count—Number of ingress flow drops due to maximum number of ingress flow drops being exceeded. • Egress Drop Flow limit drops count—Number of egress flow drops due to maximum number of egress flow drops being exceeded. |

Sample Output

show services stateful-firewall statistics extensive

```
user@host> show services stateful-firewall statistics extensive
Interface: ms-1/3/0
Service set: interface-svc-set
New flows:
  Rule Accepts: 907, Rule Discards: 0, Rule Rejects: 0
Existing flow types packet counters:
  Accepts: 3535, Drop: 0, Rejects: 0
Haripinning counters:
  Slow Path Hairpinned Packets: 0, Fast Path Hairpinned Packets: 0
Drops:
  IP option: 0, TCP SYN defense: 0
  NAT ports exhausted: 0, Sessions dropped due to subscriber flow limit: 0
Errors:
  IP: 0, TCP: 0
  UDP: 0, ICMP: 0
  Non-IP packets: 0, ALG: 0
IP errors:
  IP packet length inconsistencies: 0
  Minimum IP header length check failures: 0
  Reassembled packet exceeds maximum IP length: 0
  Illegal source address: 0
  Illegal destination address: 0
  TTL zero errors: 0, Illegal IP protocol number (0 or 255): 0
  Land attack: 0
  Non-IPv4 packets: 0, Bad checksum: 0
  Illegal IP fragment length: 0
  IP fragment overlap: 0
  IP fragment reassembly timeout: 0
  IP fragment limit exceeded: 0
  Unknown: 0
TCP errors:
  TCP header length inconsistencies: 0
  Source or destination port number is zero: 0
  Illegal sequence number and flags combination: 0
  SYN attack (multiple SYN messages seen for the same flow): 0
  First packet not a SYN message: 0
  TCP port scan (TCP handshake, RST seen from server for SYN): 0
  Bad SYN cookie response: 0
  TCP reconstructor sequence number error: 0
  TCP reconstructor retransmissions: 0
  TCP partially opened connection timeout (SYN): 0
  TCP partially opened connection timeout (SYN-ACK): 0
  TCP partially closed connection reuse: 0
  TCP 3-way error - client sent SYN+ACK: 0
  TCP 3-way error - server sent ACK: 0
  TCP 3-way error - SYN seq number retransmission mismatch: 0
  TCP 3-way error - RST seq number mismatch: 0
  TCP 3-way error - FIN received: 0
  TCP 3-way error - invalid flags (PSH, URG, ECE, CWR): 0
  TCP 3-way error - SYN recvd but no client flows: 0
  TCP 3-way error - first packet SYN+ACK: 0
  TCP 3-way error - first packet FIN+ACK: 0
  TCP 3-way error - first packet FIN: 0
  TCP 3-way error - first packet RST: 0
  TCP 3-way error - first packet ACK: 0
  TCP 3-way error - first packet invalid flags (PSH, URG, ECE, CWR): 0
  TCP Close error - no final ACK: 0
```

```
TCP Resumed Flow: 0
UDP errors:
  IP data length less than minimum UDP header length (8 bytes): 0
  Source or destination port is zero: 0
  UDP port scan (ICMP error seen for UDP flow): 0
ICMP errors:
  IP data length less than minimum ICMP header length (8 bytes): 0
  ICMP error length inconsistencies: 0
  Duplicate ping sequence number: 0
  Mismatched ping sequence number: 0
  No matching flow: 0
ALG errors:
  BOOTP: 0, DCE-RPC: 0, DCE-RPC portmap: 0
  DNS: 0, Exec: 0, FTP: 0
  H323: 0, ICMP: 0, IIOP: 0
  Login: 0, NetBIOS: 0, Netshow: 0
  Real Audio: 0, RPC: 0, RPC portmap: 0
  RTSP: 0, Shell: 0, SIP: 0
  SNMP: 0, SQLNet: 0, TFTP: 0
  Traceroute: 0
Drop Flows:
  Maximum Ingress Drop flows allowed: 20
  Maximum Egress Drop flows allowed: 20
  Current Ingress Drop flows: 0
  Current Egress Drop flows: 0
  Ingress Drop Flow limit drops count: 0
  Egress Drop Flow limit drops count: 0

**If max-drop-flows is not configured, the following is shown**
Drop Flows:
  Maximum Ingress Drop flows allowed: Default
  Maximum Egress Drop flows allowed: Default
```

show services stateful-firewall statistics application-protocol sip

| | |
|---------------------------------|--|
| Syntax | show services stateful-firewall application-protocol sip |
| Release Information | Command introduced in Junos OS Release 7.4. |
| Description | Display stateful firewall Session Initiation Protocol (SIP) statistics. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show services stateful-firewall statistics application-protocol-sip on page 2077 |
| Output Fields | Table 107 on page 2076 lists the output fields for the show services stateful-firewall statistics application-protocol-sip command. Output fields are listed in the approximate order in which they appear. |

Table 107: show services stateful-firewall statistics application-protocol-sip Output Fields

| Field Name | Field Description |
|-------------------------------|---|
| Interface | Name of an adaptive services interface. |
| Service set | Name of the service set flow. |
| ALG | Name of the application-layer gateway. |
| Active SIP call count | Number of active SIP calls. |
| Active SIP registration count | Number of active SIP registrations. |
| REGISTER | Number of new, invalid, and retransmitted register requests sent to the SIP registrar. |
| INVITE | Number of new, invalid, and retransmitted invite messages sent by user agent clients. |
| ReINVITE | Number of new, invalid, and retransmitted reinvite messages sent by user agent clients. |
| ACK | Number of new, invalid, and retransmitted ACK messages received (in response to a SIP Call Invite message). |
| BYE | Number of new, invalid, and retransmitted requests to terminate SIP dialogues. |
| CANCEL | Number of new, invalid, and retransmitted SIP request cancellations. |
| SUBSCRIBE | Number of new, invalid, and retransmitted SIP requests to subscribe for event notifications. |
| NOTIFY | Number of new, invalid, and retransmitted event notifications in SIP dialogues. |

Table 107: show services stateful-firewall statistics application-protocol-sip
Output Fields (*continued*)

| Field Name | Field Description |
|------------------------------------|---|
| OPTIONS | Number of new, invalid, and retransmitted requests to query SIP capabilities. |
| INFO | Number of new, invalid, and retransmitted requests carrying application-level information. |
| UPDATE | Number of new, invalid, and retransmitted SIP dialogue updates. |
| REFER | Number of new, invalid, and retransmitted requests to the recipient to contact a third party. |
| Provisional responses | Number of new, invalid, and retransmitted responses from the user agent server to indicate the progress of a SIP transaction. |
| OK responses to INVITES | OK responses sent from the user agent clients to user agent servers in response to Invite messages. The server can then return an ACK message. |
| OK responses to non-INVITES | OK responses to SIP messages other than an Invite message. |
| Redirection responses | Responses from the user agent server to a user agent client requesting the client to contact a different SIP uniform resource identifier (URI). |
| Request failure responses | Responses that indicate a definite failure from a particular server. The client must not retry the same request without modification after receiving this response. |
| Server failure responses | Responses that indicate a server failure. |
| Global failure responses | Responses that indicate a server has definitive information about a particular user, not just the particular instance indicated in the Request URI. |
| Invalid responses | Responses that are invalid. |
| Response (all) retransmits | Retransmissions of all responses. |
| Parser | Syntax errors, content errors, and unknown methods counted by the message parser. |

Sample Output

show services stateful-firewall statistics application-protocol-sip

```

user@host> show services stateful-firewall statistics application-protocol sip
Interface: sp-0/3/0
Service set: test_sip_777, ALG: SIP
Active SIP call count: 0, Active SIP registration count: 1

```

| | New | Invalid | Retransmit |
|----------|-----|---------|------------|
| REGISTER | 2 | | |
| INVITE | 1 | | 0 |
| ReINVITE | 1 | | |
| ACK | 1 | 0 | 0 |
| BYE | 0 | 0 | |

```
CANCEL          0          0
SUBSCRIBE       0          0
NOTIFY          0          0
OPTIONS         0          0
INFO            0          0
UPDATE          0          0
REFER           0          0
Provisional responses (18x): 1, OK responses to INVITEs: 2
OK responses to non-INVITEs: 2, Redirection (3xx) responses: 0
Request failure (4xx) responses: 0, Server failure (5xx) responses: 0
Global failure (6xx) responses: 0, Invalid responses: 0
Response (all) retransmits: 0
Parser:
  Syntax errors: 0, Content errors: 0, Unknown methods: 0
Service set: test_sip_888, ALG: SIP
Active SIP call count: 0, Active SIP registration count: 1
      New      Invalid      Retransmit
REGISTER      2
INVITE        0          0
ReINVITE      0          0
ACK           0          0
BYE           0          0
CANCEL        0          0
SUBSCRIBE     0          0
NOTIFY        0          0
OPTIONS       0          0
INFO          0          0
UPDATE        0          0
REFER         0          0
Provisional responses (18x): 0, OK responses to INVITEs: 0
OK responses to non-INVITEs: 2, Redirection (3xx) responses: 0
Request failure (4xx) responses: 0, Server failure (5xx) responses: 0
Global failure (6xx) responses: 0, Invalid responses: 0
Response (all) retransmits: 0
Parser:
  Syntax errors: 0, Content errors: 0, Unknown methods: 0
```


show services stateful-firewall subscriber-analysis

| | |
|---------------------------------|--|
| Syntax | show services stateful-firewall subscriber analysis
<interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 11.4. |
| Description | Display information about the number of active subscribers on the service physical interface card (PIC). |
| Options | none —Display standard information about all active subscribers on the PIC.

interface <i>interface-name</i> —(Optional) Display information about a particular interface. |
| Required Privilege Level | view |
| List of Sample Output | show services stateful-firewall subscriber analysis on page 2080
show services stateful-firewall subscriber-analysis on page 2080 |
| Output Fields | Table 108 on page 2079 lists the output fields for the show services stateful-firewall subscriber analysis command. Output fields are listed in the approximate order in which they appear. |

Table 108: show services stateful-firewall subscriber-analysis Output Fields

| Field Name | Field Description |
|-------------------------------------|---|
| Services PIC Name | Name of an adaptive services interface. |
| Total Subscribers Active | Total number of subscribers currently active on the service PIC. |
| Created Subscribers per Second | Rate at which subscribers are currently being created on the service PIC. |
| Deleted Subscribers per Second | Rate at which subscribers are currently being deleted on the service PIC. |
| Peak Total Subscribers Active | Highest number of subscribers that were active during the lifetime of the service PIC. |
| Peak Created Subscribers per Second | Highest rate at which subscribers were being created during the lifetime of the service PIC. |
| Peak Deleted Subscribers per Second | Highest rate at which subscribers were being deleted during the lifetime of the service PIC. |
| Number of Samples | The current sampling period lifetime. |
| Subscriber Operation: Creation | Number of sampling intervals during which a number of subscribers in the indicated range were created during the current sampling period. |
| Subscriber Operation: Deletion | Number of sampling intervals during which a number of subscribers in the indicated range were deleted during the current sampling period. |

Sample Output

show services stateful-firewall subscriber analysis

```
user@host> show services stateful-firewall subscriber analysis
```

```
Services PIC Name:    sp-2/0/0
```

```
Subscriber Analysis Statistics:
```

```
Total Subscribers Active      :100000
Created Subscribers per Second :0
Deleted Subscribers per Second :0
Peak Total Subscribers Active  :100000
Peak Created Subscribers per Second :2389
Peak Deleted Subscribers per Second :0
```

```
Subscriber Rate Data:
```

```
Number of Samples: 55
```

```
Subscriber Rate Distribution(sec)
```

```
Subscriber Operation :Creation
```

```
300000+      :0
250000 - 300000 :0
200000 - 250000 :0
160000 - 200000 :0
150000 - 160000 :0
50000 - 150000 :0
40000 - 50000 :0
30000 - 40000 :0
20000 - 30000 :0
10000 - 20000 :0
1000 - 10000 :42
0 - 1000 :1
```

```
Subscriber Operation :Deletion
```

```
300000+      :0
250000 - 300000 :0
200000 - 250000 :0
160000 - 200000 :0
150000 - 160000 :0
50000 - 150000 :0
40000 - 50000 :0
30000 - 40000 :0
20000 - 30000 :0
```

show services stateful-firewall subscriber-analysis

```
user@host> show services stateful-firewall subscriber analysis
```

```
Services PIC Name:    sp-2/0/0
```

```
Subscriber Analysis Statistics:
```

```
Total Subscribers Active      :23547
Created Subscribers per Second :2389
Deleted Subscribers per Second :0
Peak Total Subscribers Active  :23547
Peak Created Subscribers per Second :2389
Peak Deleted Subscribers per Second :0
```

```
Subscriber Rate Data:
```

```
Number of Samples: 16
```

```
Subscriber Rate Distribution(sec)
```

Subscriber Operation :Creation

```

300000+      :0
250000 - 300000 :0
200000 - 250000 :0
160000 - 200000 :0
150000 - 160000 :0
50000 - 150000 :0
40000 - 50000 :0
30000 - 40000 :0
20000 - 30000 :0
10000 - 20000 :0
1000 - 10000 :9
0 - 1000 :1

```

Subscriber Operation :Deletion

```

300000+      :0
250000 - 300000 :0
200000 - 250000 :0
160000 - 200000 :0
150000 - 160000 :0
50000 - 150000 :0
40000 - 50000 :0
30000 - 40000 :0
20000 - 30000 :0
10000 - 20000 :0
1000 - 10000 :0
0 - 1000 :0

```

Application Aware Services Operational Commands

- `clear services application-aware-access-list statistics`
- `clear services application-identification application-system-cache`
- `clear services application-identification counter`
- `clear services flows ip-action`
- `clear services local-policy-decision-function statistics`
- `request services application-identification application`
- `request services application-identification group`
- `show services application-aware-access-list flows`
- `show services application-identification application-system-cache`
- `show services application-identification counter`
- `show services application-identification group`
- `show services application-aware-access-list statistics`
- `show services application-identification application`
- `show services application-identification version`
- `show services flows`

- [show services local-policy-decision-function flows](#)
- [show services local-policy-decision-function statistics](#)

clear services application-aware-access-list statistics

| | |
|---------------------------------|---|
| Syntax | clear services application-aware-access-list statistics |
| Release Information | Command introduced in Junos OS Release 9.5. |
| Description | Clear application aware access list (AACL) statistics. |
| Options | This command has no options. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show services application-aware-access-list statistics on page 2101 |

[clear services application-identification application-system-cache](#)

| | |
|---------------------------------|--|
| Syntax | clear services application-identification application-system-cache |
| Release Information | Command introduced in Junos OS Release 9.5. |
| Description | Clear entries from application system cache. |
| Options | This command has no options. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show services application-identification application-system-cache on page 2094 |

clear services application-identification counter

| | |
|---------------------------------|---|
| Syntax | clear services application-identification counter |
| Release Information | Command introduced in Junos OS Release 9.5. |
| Description | Clear application identification counters. |
| Options | This command has no options. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show services application-identification counter on page 2096 |

clear services flows ip-action

| | |
|---------------------------------|--|
| Syntax | clear services flows ip-action |
| Release Information | Command introduced in Junos OS Release 10.0. |
| Description | Clear ip-action entries generated by the router to log, drop, or block traffic based on previous matches. The IP action options and targets are configured at the [edit security idp idp-policy <i>policy-name</i> rulebase-ips rule <i>rule-name</i> then] hierarchy level. |
| Options | This command has no options. |
| Required Privilege Level | clear |
| Output Fields | When you issue this command, you are provided feedback on the status of your request. |


Sample Output

```
user@host> clear services flows ip-action
Interface  Service set      Flows removed
ms-4/0/0   idp-service       1
```


clear services local-policy-decision-function statistics

| | |
|---------------------------------|--|
| Syntax | clear services local-policy-decision-function statistics |
| Release Information | Command introduced in Junos OS Release 9.5. |
| Description | Clear local policy decision function (L-PDF) statistics. |
| Options | This command has no options. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show services local-policy-decision-function statistics on page 2116 |

request services application-identification application

| | |
|---------------------------------|---|
| Syntax | <pre>request services application-identification application copy <i>predefined-application-name</i> request services application-identification application [disable enable] <i>predefined-application-name</i> <no-commit></pre> |
| Release Information | Command introduced in Release 11.4 of Junos OS. |
| Description | Copy, disable, or enable a predefined application signature. |
| Options | <p>copy—(Optional) Copy a predefined application signature from the database to the configuration and change the name (for example, my:FTP). The ID and order will be generated automatically. Do not name your custom application signature with the “junos” prefix; this prefix is reserved for predefined application signatures. You can copy the same predefined application signature only once; duplicate custom signatures are not allowed. The copy command does not initiate signature recompilation.</p> <hr/> <p> NOTE: In configuration mode, if an uncommitted action is pending, the request services application-identification application copy command will fail. Uncompiled application signatures are shown as uncommitted in the show services application-identification application [summary detail] command.</p> <hr/> <p>disable—(Optional) Disable a predefined application signature, initiate signature recompilation, and commit all pending uncompiled signatures to the configuration. Include the <no-commit> keyword to defer signature recompilation.</p> <p>The following conditions apply:</p> <ul style="list-style-type: none"> • You cannot disable a predefined application signature that is referenced by an active security policy or custom application signature. First modify or deactivate the policy or custom application signature. • If you disable an application signature, for example, junos:HTTP, that has nested applications, the nested applications will not be recognized. <p>enable—(Optional) Enable a predefined application signature, initiate signature recompilation, and commit all pending uncompiled signatures to the configuration. Include the <no-commit> keyword to defer signature recompilation.</p> <p>no-commit—(Optional) Enables you to enter multiple enable and disable commands before initiating signature recompilation (which takes some time) and committing the configuration. Uncompiled application signatures are shown as uncommitted in the show services application-identification application [summary detail] command.</p> |
| Required Privilege Level | maintenance |

Related Documentation

- [show services application-identification application on page 2103](#)

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

[request services application-identification application copy](#)

```
user@host> request services application-identification application junos:FTP copy
application package 63 copied successfully.
```



[request services application-identification application disable](#)

```
user@host> request services application-identification application disable junos:163
Please wait while we are re-compiling signatures ...
Please wait while we are re-compiling signatures ...
Please wait while we are re-compiling signatures ...
Please wait while we are re-compiling signatures ...
Disable application junos:163 succeed.
```

[request services application-identification application disable no-commit](#)

```
user@host> request services application-identification application disable
junos:FACEBOOK-SOCIALRSS no-commit
Disable application junos:FACEBOOK-SOCIALRSS succeed. It is not committed yet.
```

request services application-identification group

| | |
|---------------------------------|---|
| Syntax | <code>request services application-identification group [copy disable enable]
predefined-application-group-name</code> |
| Release Information | Command introduced in Release 11.4 of Junos OS. |
| Description | Copy, disable, or enable a predefined application signature group. |
| Options | <p>copy—(Optional) Copy a predefined application signature group from the database to the configuration and change the name (for example, my:FTP). The ID and order will be generated automatically. Do not name your custom application signature group with the “junos” prefix; this prefix is reserved for predefined application signature groups. You can copy the same predefined application signature group only once; duplicate custom signature groups are not allowed.</p> <p>.....</p> <p> NOTE: In configuration mode, if an uncommitted action is pending, the <code>request services application-identification group copy</code> command will fail.</p> <p>.....</p> <p>disable—(Optional) Disable a predefined application signature group.</p> <p>.....</p> <p> NOTE: You cannot disable a predefined application signature group that is referenced by an active security policy or custom application signature group. First modify or deactivate the policy or custom application signature group.</p> <p>.....</p> <p>enable—(Optional) Enable a predefined application signature group.</p> |
| Required Privilege Level | maintenance |
| Related Documentation | <ul style="list-style-type: none">• show services application-identification group on page 2099 |
| Output Fields | When you enter this command, the system provides feedback on the status of your request. |

Sample Output

request services application-identification group

```
user@host> request services application-identification group copy junos:SYBASE
group 1040 copied successfully.
```

show services application-aware-access-list flows

| | |
|---------------------------------|--|
| Syntax | show services application-aware-access-list flows
<interface <i>interface-name</i>>
<subscriber <i>subscriber-name</i>> |
| Release Information | Command introduced in Junos OS Release 10.1.
Offload status for flows using Juniper Forwarding Mechanism (JFM) added in Junos OS Release 12.1. |
| Description | Display application-aware-access-list (AACL) flows. Offloading using JFM is supported only on MX Series routers with Modular Port Concentrators (MPCs). |
| Options | interface <i>interface-name</i> —Displays AACL flows for the specified interface(s) only. The keyword, interface, must be appended to the command.

subscriber <i>subscriber-name</i> —Displays AACL flows for the specified subscriber(s) only. The keyword, subscriber, must be appended to the command. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> <i>Application Aware Services Interfaces Feature Guide for Routing Devices</i> |
| List of Sample Output | show services application-aware-access-list flows by interface on page 2092
show services application-aware-access-list flows by subscriber on page 2092
show services application-aware-access-list flows by subscriber for offloading using JFM on page 2093 |
| Output Fields | Table 109 on page 2091 lists the output fields for the show services application-aware-access-list flows command. Output fields are listed in the approximate order in which they appear. |

Table 109: show services application-aware-access-list flows Output Fields

| Field Name | Field Description | Level of Output |
|----------------|--|-----------------|
| 5-tuple | This field comprises five components of the given flow. The components are: <ul style="list-style-type: none"> • Src IP • Dest IP • Src Port • Dest Port • Protocol | All levels |
| Application-ID | The identification number associated with the application. | All levels |
| Dir | The direction in terms of input or output. <ul style="list-style-type: none"> • Input (I) • Output (O) | All levels |

Table 109: show services application-aware-access-list flows Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------|--|-----------------|
| Off | The status of offload to Packet Forwarding Engine. The various options are: <ul style="list-style-type: none"> • Not Offloaded (-) • Policer Offloaded, Flow Not Offloaded (P) • Policer Not Offloaded, Flow Offloaded (F) • Policer and Offloaded (P+F) | All levels |
| Off | The status of offload to Packet Forwarding Engine using JFM. The various options are: <ul style="list-style-type: none"> • Not Offloaded (-) • Offload requested but not completed (R) • Offload requested and completed (O) | All levels |
| Actions | The types of actions displayed are: <ul style="list-style-type: none"> • discard: (D) • accept : A • accept, count [T]: C-A or C-G or C-T • accept, fwd-class [C]: FC • accept, policer [P]: P • accept, count [T], fwd-class [C]: C-T+FC • accept, count [T], policer [P]: C-T+P • accept, fwd-class [C], policer [P]: FC+P • accept, count[T],fwd-class[C],policer[P]: C-T+FC+P | All levels |

Sample Output

show services application-aware-access-list flows by interface

```

user@host>show services application-aware-access-list flows interface ge-1/0/5.0
Interface: ge-1/0/5.0
service-set: aac1-countApps
service-set interface: ms-0/0/0
Currently active flows: 2
High watermark flows: 2

5-tuple                                     Application-ID
Dir Off Action
-----
--- ---
      1.0.5.2:47072-> 10.10.254.116:80 ,6 junos:http [64]
I  - C-T
      10.10.254.116:80 ->      1.0.5.2:47072,6 junos:http [64]
O  - C-T

```

show services application-aware-access-list flows by subscriber

```

user@host>show services application-aware-access-list flows subscriber user@juniper.net
Subscriber: user@juniper.net

Service-set: ss1

```

```
Service-set interface: ms-2/0/0
Currently active flows: 4
High watermark flows: 40
```

| 5-tuple | Application-ID | Dir | Off | Action |
|--|-----------------|-----|-----|----------|
| 150.100.100.100:20109->160.200.200.200:80,17 | junos:http [64] | I | - | C-T+FC+P |
| 160.200.200.200:80->150.100.100.100:20109,17 | junos:http [64] | O | - | C-T+FC+P |
| 150.100.100.100:20108->160.100.100.100:80,17 | junos:http [64] | I | P+F | C-T+FC+P |
| 160.100.100.100:80->150.100.100.100:20108,17 | junos:http [64] | O | P+F | C-T+FC+P |

show services application-aware-access-list flows by subscriber for offloading using JFM

```
user@host>show services application-aware-access-list flows subscriber user@juniper.net
Subscriber: user@juniper.net
```

```
Service-set: ssl
Service-set interface: ms-2/0/0
Currently active flows: 4
High watermark flows: 40
```

| 5-tuple | Application-ID | Dir | Off | Action |
|---|-----------------|-----|-----|--------|
| 150.100.100.100:20109->160.200.200.200:80,17 | junos:http [64] | | | I |
| - C-T+FC+P | | | | |
| 160.200.200.200:80 ->150.100.100.100:20109,17 | junos:http [64] | | | O |
| - C-T+FC+P | | | | |
| 150.100.100.100:20108->160.100.100.100:80,17 | junos:http [64] | | | I |
| R C-T+FC+P | | | | |
| 160.100.100.100:80 ->150.100.100.100:20108,17 | junos:http [64] | | | O |
| O C-T+FC+P | | | | |

show services application-identification application-system-cache

Syntax `show application-identification application-system-cache
<interface interface-name>`

Release Information Command introduced in Junos OS Release 9.5.
interface option added in Junos OS Release 10.1.

Description Display the database of cached values stored by the application identification (APPID) system.



NOTE: The `show services application-identification application-system-cache` command gives the information only when the application identifier (AI) is matched with the signature.

Options `interface interface-name`—Displays the services interfaces to query.

Required Privilege Level view

List of Sample Output [show application-identification application-system-cache on page 2094](#)

Output Fields [Table 110 on page 2094](#) lists the output fields for the `command-name` command. Output fields are listed in the approximate order in which they appear.

Table 110: show application-identification application-system-cache Output Fields

| Field Name | Field Description | Level of Output |
|-------------|---------------------|-----------------|
| IP address | IP address. | All levels |
| Port | Port number. | All levels |
| Protocol | Protocol name. | All levels |
| Application | Application number. | All levels |
| CPU | CPU number | All levels |

Sample Output

show application-identification application-system-cache

```
user@host> show application-identification application-system-cache interface ms-1/0/0
pic: 2/0
```

```
IP address      Port      Protocol  Application  CPU
```


10.1.1.2

81

TCP

63

18

show services application-identification counter

| | |
|---------------------------------|---|
| Syntax | show services application-identification counter
<interface <i>interface-name</i>> |
| Release Information | Command introduced in Junos OS Release 9.5.
interface option added in Junos OS Release 10.1. |
| Description | Display application identification (APPID) counter statistics. |
| Options | interface <i>interface-name</i> —Displays the services interfaces to query. |
| Required Privilege Level | view |
| List of Sample Output | show services application-identification counter on page 2097
show services application-identification counter on page 2097 |
| Output Fields | Table 111 on page 2096 lists the output fields for the show services application-identification counter command. Output fields are listed in the approximate order in which they appear. |

Table 111: show services application-identification counter Output Fields

| Field Name | Field Description |
|--|---|
| pic | PIC number. |
| Total sessions | Total number of sessions. |
| Total identified sessions | Total number of identified sessions. |
| Total unidentified sessions | Total number of unidentified sessions. |
| Total identified-by-address sessions | Number of sessions identified by address. |
| Total unidentified-by-address sessions | Number of sessions not identified by address. |
| Total identified-by-port sessions | Number of sessions identified by port. |
| Total unidentified-by-port sessions | Number of sessions not identified by port. |
| Total identified-by-icmp sessions | Number of sessions identified by ICMP. |
| Total unidentified-by-icmp sessions | Number of sessions not identified by ICMP. |
| Total identified-by-ip-protocol sessions | Number of sessions identified by IP protocol. |
| Total unidentified-by-ip-protocol sessions | Number of sessions not identified by IP protocol. |
| Total identified-by-signature sessions | Number of sessions identified by signature. |

Table 111: show services application-identification counter Output Fields (*continued*)

| Field Name | Field Description |
|--|---|
| Total unidentified-by-signature sessions | Number of sessions not identified by signature. |
| Total unspecified encrypted sessions | Number of encrypted sessions not specified by normal processes. |
| Total encrypted P2P sessions | Number of encrypted point-to-point sessions. |
| Total application system cache hits | Number of sessions found in the application system cache. |
| Total application system cache misses | Number of sessions not found in the application system cache. |
| Total identified-by-protocol sessions | Number of sessions identified by protocol. |
| Total unidentified-by-protocol sessions | Number of sessions not identified by protocol. |

Sample Output

show services application-identification counter

```

user@host> show services application-identification counter interface ms-1/0/0
Counter Statistics:
  pic: 1/1
  Total sessions: 11
  Total identified sessions: 11
  Total un-identified sessions: 0
Address Method
  Total identified-by-address sessions: 0
  Total unidentified-by-address sessions: 11
Port Method
  Total identified-by-port sessions: 1
  Total unidentified-by-port sessions: 0
  Total identified-by-icmp sessions: 0
  Total unidentified-by-icmp sessions: 0
  Total identified-by-ip-protocol sessions: 0
  Total unidentified-by-ip-protocol sessions: 0
Signature Method
  Total identified-by-signature sessions: 11
  Total unidentified-by-signature sessions: 0
  Total unspecified encrypted sessions: 2
  Total encrypted P2P sessions: 2
  Total application system cache hits: 10
  Total application system cache misses: 1
Protocol Method
  Total identified-by-protocol sessions: 0
  Total unidentified-by-protocol sessions: 0

```

show services application-identification counter

```

user@host> show services application-identification counter interface ams0
Counter Statistics:
  pic: ams0
  Total sessions: 20
  Total identified sessions: 20
  Total un-identified sessions: 0

```

```
Protocol Method
  Total identified-by-protocol sessions: 0
  Total un-identified-by-protocol sessions: 0
Address Method
  Total identified-by-address sessions: 0
  Total un-identified-by-address sessions: 0
Port Method
  Total identified-by-port sessions: 0
  Total un-identified-by-port sessions: 0
  Total identified-by-icmp sessions: 0
  Total un-identified-by-icmp sessions: 0
  Total identified-by-ip-protocol sessions: 0
  Total un-identified-by-ip-protocol sessions: 0
Signature Method
  Total identified-by-signature sessions: 20
  Total identified-by-signature uni-directional sessions: 0
  Total un-identified-by-signature sessions: 0
  Total application system cache hits: 0
  Total application system cache misses: 0
```

show services application-identification group

| | |
|---------------------------------|---|
| Syntax | <code>show services application-identification group [detail <i>application-group name</i>] summary]</code> |
| Release Information | Command introduced in Release 11.4 of Junos OS. |
| Description | Display detailed or summary information about a specified application signature group or all application signature groups. Both custom and predefined application signature groups can be displayed. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • request services application-identification group on page 2090 |
| List of Sample Output | show services application-identification group summary on page 2099
show services application-identification group detail on page 2100 |
| Output Fields | Table 112 on page 2099 lists the output fields for the show services application-identification group command. Output fields are listed in the approximate order in which they appear. |

Table 112: show services application-identification group Output Fields

| Field Name | Field Description |
|-----------------------------|--|
| Description | Description of the specified application in the detailed display. |
| Group ID or ID | The unique ID number of an application signature or application signature group. ID numbers 1 through 32,767 are automatically generated for predefined application signatures and application signature groups; these IDs do not change. ID numbers for custom application signatures and application signature groups use ID numbers 32,768 to 65,534. |
| Disabled | The status of the application signature group and whether the signature method is currently used to identify this application. The default is No. |
| Application Group(s) | The application signature groups present. |
| Applications | The application signatures associated with this application signature group. |

Sample Output

show services application-identification group summary

```

user@host> show services application-identification group summary
Application Group(s): 24
Application Groups
my:enterprise                Disabled  ID
                             No             32770
junos:enterprise:voip        No             25
junos:peer-to-peer:voip      No             24
junos:peer-to-peer:chat      No             23
junos:peer-to-peer:file-sharing No             22
...
```

show services application-identification group detail

```
user@host> show services application-identification group detail junos:social-networking
Description: Detection for social networking sites such as Myspace, Facebook and
similar.
Group ID: 8
Disabled: no
Application-groups:
    junos:social-networking:facebook;
    junos:social-networking:myspace;

Applications:
    junos:4CHAN;
    junos:ADULTFRIENDFINDER;
    junos:BAD00;
    junos:BEBO;
    junos:BLOGGER-POST;
    junos:BLOGSPOT-POST;
```

show services application-aware-access-list statistics

| | |
|---------------------------------|---|
| Syntax | show services application-aware-access-list statistics
<interface <i>interface-name</i>>
<subscriber <i>subscriber-name</i>> |
| Release Information | Command introduced in Junos OS Release 9.5. |
| Description | Display application-aware-access-list (AACL) statistics. |
| Options | interface <i>interface-name</i> —(Optional) Displays AACL statistics for the specified interface(s) only.

subscriber <i>subscriber-name</i> —(Optional) Displays AACL statistics for the specified subscriber(s) only. |
| Required Privilege Level | view |
| List of Sample Output | show services application-aware-access-list statistics by interface on page 2102
show services application-aware-access-list statistics by subscriber on page 2102 |
| Output Fields | Table 113 on page 2101 lists the output fields for the show services application-aware-access-list statistics command. Output fields are listed in the approximate order in which they appear. |

Table 113: show services application-aware-access-list statistics Output Fields

| Field Name | Field Description | Level of Output |
|------------------------------|-------------------------------|-------------------|
| Interface | Interface name. | Subscriber option |
| Subscriber | Subscriber identifier. | Interface option |
| Service-set-interface | Service set interface name. | All levels |
| Service set | Service set name. | All levels |
| Application group | Application group identifier. | All levels |
| Packets in | Number of ingress packets. | All levels |
| Bytes in | Number of ingress bytes. | All levels |
| Packets out | Number of egress packets. | All levels |
| Bytes out | Number of egress bytes. | All levels |

Sample Output

show services
application-aware-access-list
statistics by interface

```
user@host> show services application-aware-access-list statistics interface ge-0/0/0.100
Subscriber: user@juniper.net
```

```
service-set: IDP
service-set interface: ms-2/0/0
```

| Application group | Application | Packets in | Bytes in |
|-------------------|----------------|------------|----------|
| Packets out | Bytes out | | |
| | junos:ftp [63] | 5 | 334 |
| 6 | 346 | | |

show services
application-aware-access-list
statistics by subscriber

```
user@host> show services application-aware-access-list statistics subscriber user@juniper.net
Interface: ge-1/1/0.0
```

```
Service-set-interface: ms-1/3/0
Service set: aacl-svc-set
```

Application-aware-access-list statistics

| Application group | Packets in | Bytes in | Packets out | Bytes |
|-------------------|------------|----------|-------------|-------|
| out | | | | |
| P2P | | 400 | 32025 | 200 |
| | 16284 | | | |
| FTP | | 20000 | 5231000 | 100 |
| | 8700 | | | |

show services application-identification application

| | |
|---------------------------------|--|
| Syntax | show services application-identification application [detail <i>application-name</i>] summary] |
| Release Information | Command introduced in Release 11.4 of Junos OS. |
| Description | Display detailed information about a specified application signature, all application signatures, or a summary of the existing application signatures and nested application signatures. Both custom and predefined application signatures and nested application signatures can be displayed. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • request services application-identification application on page 2088 |
| List of Sample Output | show services application-identification application summary on page 2104
show services application-identification application detail on page 2104 |
| Output Fields | Table 114 on page 2103 lists the output fields for the show services application-identification application command. Output fields are listed in the approximate order in which they appear. |

Table 114: show services application-identification application Output Fields

| Field Name | Field Description |
|-----------------------|---|
| Application(s) | The number of application signatures present. |
| Nested Application(s) | The number of nested application signatures present. |
| Application Name | Name of the predefined or custom application signature. Must be a unique name with a maximum length of 32 characters. |
| Application Group | Name of the application signature group associated with this application signature/nested application signature. Must be a unique name with a maximum length of 32 characters. |
| Disabled | The status of the application signature or nested application signature and whether the signature method is currently used to identify this application. The default is No. |
| ID | The unique ID number of an application signature or nested application signature. ID numbers 1 through 32,767 are automatically generated for predefined application signatures and nested application signatures; these IDs do not change. ID numbers for custom application signatures and nested application signatures use ID numbers 32,768 to 65,534. |
| Order | A unique number used to specify priority when multiple patterns are matched for the same session. The lowest order number takes the highest priority. |

Table 114: show services application-identification application Output Fields (*continued*)

| Field Name | Field Description |
|--|---|
| Application Tags | General information about this application type, for example, associated risk factors, technology, type of traffic, and so on.

Support of application signature tags is dependent on the version of the loaded signature database. Please refer to the Juniper Networks security portal for further information. |
| Port Mapping: Default ports | The default port for this application type. |
| Signature: Port range | Default ranges: TCP/0 through 65,535; UDP/0 through 65,535 (optional). |
| Client-to-server: DFA Pattern | Pattern-matching scheme used for client-to-server traffic. Maximum length is 1023 (optional). |
| Client-to-server: Regex Pattern | A compatible regular expression used to match client-to-server traffic. |
| Server to-client: DFA Pattern | Pattern-matching scheme used for server-to-client traffic. Maximum length is 1023 (optional). |
| Server-to-client: Regex Pattern | A compatible regular expression used to match server-to-client traffic. |
| Minimum Data | The minimum number of bytes or packets to apply to the DFA pattern. Default is 10; range is 4 through 1024. |

Sample Output

show services application-identification application summary

```

user@host> show services application-identification application summary
Application(s): 150
Nested Application(s): 600

Application      Disabled      ID      Order
...             ...          ...     ...
junos:FTP        No           63      59
junos:HTTP       Yes          64      122
...             ...          ...     ...
my:APPLICATION-A No           700     730
my:APPLICATION-B No           701     731
...             ...          ...     ...

```

show services application-identification application detail

```

user@host> show services application-identification detail junos:FTP
Description: This signature detects the File Transfer Protocol (FTP), which
provides facilities for transferring files to and from remote computer systems.
It usually runs on TCP port 21.
Application ID: 63
Disabled: No
Number of Parent Group(s): 1
Application Groups:
  junos:file-server
Application Tags:
  characteristic      : Supports File Transfer

```

```

characteristic      : Known Vulnerabilities
characteristic      : Capable of Tunneling
risk               : 3
category           : FILE-SERVER
Port Mapping:
  Default ports: TCP/21
Signature:
  Port range: TCP/0-24,26-65535
  Client-to-server
    DFA Pattern:
      \[(USER|STAT|PORT|CHMOD|ACCOUNT|BYE|ASCII|GLOB|HELP|AUTH|SYST|QUIT|STOR|PASV|QWD|PWD|MDTM|FEAT|OPTS)\](\s|\x0d
      0a\x|\x0a\x).*
    Regex Pattern: None
  Server-to-client
    DFA Pattern: (220|230|331|530)[\s\-.]*
    Regex Pattern: None
  Minimum data client-to-server: 8
  Minimum data server-to-client: 8
  Order: 71

```

show services application-identification version

| | |
|---------------------------------|---|
| Syntax | show services application-identification version |
| Release Information | Command introduced in Release 10.2 of Junos OS. |
| Description | Display the Junos OS application package version. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• <i>request services application-identification download</i> |
| List of Sample Output | show services application-identification version on page 2106 |

Sample Output

show services application-identification version

The following output shows that the application package version is 1608.

```
user@host> show services application-identification version
Application package version: 1608
```

show services flows

Syntax show services flows
 <all | brief | extensive | terse>
 <application-protocol *protocol*>
 <count>
 <destination-port *destination-port*>
 <destination-prefix *destination-prefix*>
 <interface *interface-name*>
 <limit *number*>
 <protocol *protocol*>
 <service-set *service-set*>
 <source-port *source-port*>
 <source-prefix *source-prefix*>

Release Information Command introduced in Junos OS Release 9.5.
all option introduced in Junos OS Release 11.1.
application-protocol option introduced in Junos OS Release 11.1.

Description Display flow session table entries.

Options **none**—Display standard information about all flows.

all | brief | extensive | terse—(Optional) Display the specified level of output.

application-protocol—(Optional) Display information about one of the following application protocols:

- **bootp**—Bootstrap protocol
- **dce-rpc**—Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—Domain Name System protocol
- **exec**—Exec
- **ftp**—File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol
- **iiop**—Internet Inter-ORB Protocol
- **login**—Login
- **netbios**—NetBIOS
- **netshow**—NetShow
- **pptp**—Point-to-Point Tunneling Protocol
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol

- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **shell**—Shell
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet
- **talk**—Talk Program
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame



NOTE: The flows for the DCE RPC ALG match the flows for the DCE RPC Portmap ALG. The flows for the RPC ALG match the flows for the RPC Portmap ALG.

count—(Optional) Display a count of the matching entries.

destination-port *destination-port*—(Optional) Display information for a particular destination port. The range of values is from 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Display information for a particular destination prefix.

interface *interface-name*—(Optional) Display information about a particular interface. On M Series and T Series routers, *interface-name* can be **ms-fpc/pic/port** or **rspnumber**.

limit *number*—(Optional) Maximum number of entries to display.

protocol *protocol*—(Optional) Display information about one of the following IP types:

- **number**—Numeric protocol value from 0 to 255
- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **icmp6**—Internet Control Message Protocol version 6
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol

- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Transmission Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Display information for a particular service set.

source-port *source-port*—(Optional) Display information for a particular source port. The range of values is from 0 to 65535.

source-prefix *source-prefix*—(Optional) Display information for a particular source prefix.

Required Privilege Level view

Related Documentation

- *clear services flows*

List of Sample Output [show services flows on page 2110](#)
[show services flows all on page 2110](#)
[show services flows brief on page 2111](#)
[show services flows extensive on page 2111](#)
[show services flows application-protocol on page 2111](#)
[show services flows count on page 2111](#)
[show services flows destination port on page 2111](#)
[show services flows destination prefix on page 2112](#)
[show services flows interface on page 2112](#)
[show services flows protocol on page 2112](#)
[show services flows service-set on page 2112](#)
[show services flows source port on page 2112](#)
[show services flows source prefix on page 2112](#)

Output Fields [Table 115 on page 2109](#) lists the output fields for the **show services flows** command. Output fields are listed in the approximate order in which they appear.

Table 115: show services flows Output Fields

| Field Name | Field Description | Level of Output |
|--------------------|---|-------------------|
| Interface | Name of the interface. | All levels |
| Service set | Name of a service set. Individual empty service sets are not displayed. If no service set has any flows, a flow table header is displayed for each service set. | All levels |
| Flow Count | Number of flows in a session. | count only |

Table 115: show services flows Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------|--|------------------|
| Flow or Flow Prot | Protocol used for this flow. | All levels |
| Source | Source prefix of the flow in the format <i>source-prefix:port</i> . For ICMP flows, port information is not displayed. | All levels |
| Dest | Destination prefix of the flow. For ICMP flows, port information is not displayed. | All levels |
| State | Status of the flow: <ul style="list-style-type: none"> • Drop—Drop all packets in the flow without response. • Forward—Forward the packet in the flow without looking at it. • Reject—Drop all packets in the flow with response. • Watch—Inspect packets in the flow. | All levels |
| Dir | Direction of the flow: input (I) or output (O). | All levels |
| Frm count | Number of frames in the flow. | All levels |
| Byte count | Number of bytes in the flow. | extensive |
| Flow role | Flow role. | extensive |
| Timeout | Timeout value. | extensive |
| Flow path | Flow path: symmetric or asymmetric. | extensive |

Sample Output

show services flows

```

user@host> show services flows
Interface: ms-2/0/0, Service set: IDP
Flow
TCP      10.2.2.2:33656 -> 10.1.1.2:80    Forward I      6
TCP      10.1.1.2:80   -> 10.2.2.2:33656 Forward O      5
ICMP     10.1.1.2       -> 10.2.2.2       Forward I     102
ICMP     10.2.2.2       -> 10.1.1.2       Forward O     102
ICMP     10.2.2.2       -> 10.1.1.2       Forward I      97
ICMP     10.1.1.2       -> 10.2.2.2       Forward O      97

```

show services flows all

```

user@host> show services flows all
Interface: ms-2/0/0, Service set: idp-1
Flow
TCP      10.1.1.2:32769 -> 20.1.1.2:80    Forward I    353431
TCP      20.1.1.2:80   -> 10.1.1.2:32769 Forward O    353429
TCP      10.1.1.2:32771 -> 20.1.1.2:80    Forward I    353562
TCP      20.1.1.2:80   -> 10.1.1.2:32771 Forward O    353560

```


| | | | | | | |
|-----|----------------|----|----------------|---------|---|--------|
| TCP | 10.1.1.2:32770 | -> | 20.1.1.2:80 | Forward | I | 353577 |
| TCP | 20.1.1.2:80 | -> | 10.1.1.2:32770 | Forward | O | 353575 |
| TCP | 10.1.1.2:32768 | -> | 20.1.1.2:80 | Forward | I | 353610 |
| TCP | 20.1.1.2:80 | -> | 10.1.1.2:32768 | Forward | O | 353608 |
| TCP | 10.1.1.2:32777 | -> | 20.1.1.2:80 | Forward | I | 353625 |
| TCP | 20.1.1.2:80 | -> | 10.1.1.2:32777 | Forward | O | 353624 |
| TCP | 10.1.1.2:32776 | -> | 20.1.1.2:80 | Forward | I | 353643 |
| TCP | 20.1.1.2:80 | -> | 10.1.1.2:32776 | Forward | O | 353642 |
| TCP | 10.1.1.2:32775 | -> | 20.1.1.2:80 | Forward | I | 353658 |
| TCP | 20.1.1.2:80 | -> | 10.1.1.2:32775 | Forward | O | 353657 |
| TCP | 10.1.1.2:32774 | -> | 20.1.1.2:80 | Forward | I | 353676 |
| TCP | 20.1.1.2:80 | -> | 10.1.1.2:32774 | Forward | O | 353674 |
| TCP | 10.1.1.2:32773 | -> | 20.1.1.2:80 | Forward | I | 353692 |
| TCP | 20.1.1.2:80 | -> | 10.1.1.2:32773 | Forward | O | 353690 |
| TCP | 10.1.1.2:32772 | -> | 20.1.1.2:80 | Forward | I | 353704 |
| TCP | 20.1.1.2:80 | -> | 10.1.1.2:32772 | Forward | O | 353702 |

show services flows brief

The output for the **show services flows brief** command is identical to that for the **show services flows** command. For sample output, see [show services flows](#).

show services flows extensive

```
user@host> show services flows extensive
Interface: ms-2/0/0, Service set: IDP
Flow                                     State  Dir  Frm count
TCP      10.2.2.2:33656 ->      10.1.1.2:80    Forward I           6
  Byte count: 346
  Flow role: Unknown, Timeout: 0, Flow path: Asymmetric
TCP      10.1.1.2:80 ->      10.2.2.2:33656 Forward O           5
  Byte count: 334
  Flow role: Unknown, Timeout: 0, Flow path: Symmetric
ICMP      10.1.1.2 ->      10.2.2.2      Forward I          144
  Byte count: 12096
  Flow role: Unknown, Timeout: 0, Flow path: Symmetric
ICMP      10.2.2.2 ->      10.1.1.2      Forward O          144
  Byte count: 12096
  Flow role: Unknown, Timeout: 0, Flow path: Symmetric
```

show services flows application-protocol

```
user@router> show services flows application-protocol dce-rpc
Interface: ms-2/0/0, Service set: ss-1
Flow                                     State  Dir  Frm count
TCP      192.168.200.65:1260 -> 192.168.200.69:5315 Forward I          14
TCP      192.168.200.69:5315 ->  16.16.16.16:1031 Forward O          11
TCP      192.168.200.65:1251 -> 192.168.200.69:1026 Forward I           7
TCP      192.168.200.69:1026 ->  16.16.16.16:1029 Forward O           5
```

show services flows count

```
user@host> show services flows count
Interface  Service set  Flow count
ms-2/0/0   IDP          6
```

show services flows destination port

```
user@router> show services flows destination-port 80
```

```

Interface: ms-2/0/0, Service set: IDP
Flow      State  Dir  Frm count
TCP       10.2.2.2:33656 -> 10.1.1.2:80 Forward I      6

```

show services flows destination prefix

```

user@router> show services flows destination-prefix 10.1.1.2
Interface: ms-2/0/0, Service set: IDP
Flow      State  Dir  Frm count
TCP       10.2.2.2:33656 -> 10.1.1.2:80 Forward I      6
ICMP      10.2.2.2      -> 10.1.1.2      Forward O     137
ICMP      10.2.2.2      -> 10.1.1.2      Forward I     132

```

show services flows interface

```

user@router> show services flows interface ms-2/0/0
Interface: ms-2/0/0, Service set: IDP
Flow      State  Dir  Frm count
TCP       10.2.2.2:33656 -> 10.1.1.2:80 Forward I      6
TCP       10.1.1.2:80 -> 10.2.2.2:33656 Forward O      5
ICMP      10.1.1.2      -> 10.2.2.2      Forward I     162
ICMP      10.2.2.2      -> 10.1.1.2      Forward O     162
ICMP      10.2.2.2      -> 10.1.1.2      Forward I     157
ICMP      10.1.1.2      -> 10.2.2.2      Forward O     157

```

show services flows protocol

```

user@router> show services flows protocol icmp
Interface: ms-2/0/0, Service set: IDP
Flow      State  Dir  Frm count
ICMP      10.1.1.2      -> 10.2.2.2      Forward I     202
ICMP      10.2.2.2      -> 10.1.1.2      Forward O     202
ICMP      10.2.2.2      -> 10.1.1.2      Forward I     197
ICMP      10.1.1.2      -> 10.2.2.2      Forward O     197

```

show services flows service-set

```

user@router> show services flows service-set sample
Interface: ms-2/0/0, Service set: sample
Flow      State  Dir  Frm count
TCP       10.2.2.2:33656 -> 10.1.1.2:80 Forward I      6
TCP       10.1.1.2:80 -> 10.2.2.2:33656 Forward O      5
ICMP      10.1.1.2      -> 10.2.2.2      Forward I     220
ICMP      10.2.2.2      -> 10.1.1.2      Forward O     220
ICMP      10.2.2.2      -> 10.1.1.2      Forward I     215
ICMP      10.1.1.2      -> 10.2.2.2      Forward O     215

```

show services flows source port

```

user@router> show services flows source-port 0
Interface: ms-2/0/0, Service set: IDP
Flow      State  Dir  Frm count
TCP       10.2.2.2:33656 -> 10.1.1.2:80 Forward I      6
TCP       10.1.1.2:80 -> 10.2.2.2:33656 Forward O      5
ICMP      10.1.1.2      -> 10.2.2.2      Forward I     235
ICMP      10.2.2.2      -> 10.1.1.2      Forward O     235
ICMP      10.2.2.2      -> 10.1.1.2      Forward I     230
ICMP      10.1.1.2      -> 10.2.2.2      Forward O     230

```

show services flows source prefix

```

user@router> show services flows source-prefix 10.2.2.2

```

Interface: ms-2/0/0, Service set: IDP

| Flow | | | | State | Dir | Frm count |
|------|----------------|----|----------------|---------|-----|-----------|
| TCP | 10.2.2.2:33656 | -> | 10.1.1.2:80 | Forward | I | 6 |
| TCP | 10.1.1.2:80 | -> | 10.2.2.2:33656 | Forward | O | 5 |
| ICMP | 10.1.1.2 | -> | 10.2.2.2 | Forward | I | 235 |
| ICMP | 10.2.2.2 | -> | 10.1.1.2 | Forward | O | 235 |
| ICMP | 10.2.2.2 | -> | 10.1.1.2 | Forward | I | 230 |
| ICMP | 10.1.1.2 | -> | 10.2.2.2 | Forward | O | 230 |

show services local-policy-decision-function flows

| | |
|---------------------------------|--|
| Syntax | show services local-policy-decision-function flows (interface <i>interface-name</i> subscriber <i>subscriber-name</i>) |
| Release Information | Command introduced in Junos OS Release 9.5. |
| Description | Display local policy decision function (L-PDF) flows. |
| Options | interface <i>interface-name</i> —Display L-PDF flows for the specified interfaces only.
subscribers <i>subscriber-name</i> —Display L-PDF flows for the specified subscribers only. |
| Required Privilege Level | view |
| List of Sample Output | show services local-policy-decision-function flows by interface on page 2115
show services local-policy-decision-function flows by subscriber on page 2115 |
| Output Fields | Table 116 on page 2114 lists the output fields for the show services local-policy-decision-function flows command. Output fields are listed in the approximate order in which they appear. |

Table 116: show services local-policy-decision-function flows Output Fields

| Field Name | Field Description |
|-------------------------------|--|
| Interface | Interface name. |
| service-set | Service set name. |
| service-set-interface | Service set interface name. |
| Currently active flows | Number of currently active flows. |
| High watermark flows | Maximum number of flows. |
| Protocol | (With interface option) Protocol identifier. |
| Source address | (With interface option) Source address. |
| Source port | (With interface option) Source port. |
| Destination address | (With interface option) Destination address. |
| Destination port | (With interface option) Destination port. |
| Application | (With interface option) Application name. |
| Application group | (With interface option) Application group identifier. |

Sample Output

show services local-policy- decision-function flows by interface

```
user@host> show services local-policy-decision-function flows subscriber user@juniper.net
Interface: ge-0/0/5.26
```

```
service-set: aac1_ms30
service-set interface: ms-3/0/0
```

```
Currently active flows: 0
High watermark flows: 0
```

show services local-policy- decision-function flows by subscriber

```
user@host> show services local-policy-decision-function flows interface ge-1/1/0
Interface: ge-1/1/0.0
```

```
service-set: IDP
service-set interface: ms-2/0/0
```

```
Currently active flows: 2
High watermark flows: 2
```

| Protocol | Source address | Source port | Destination address | Destination port |
|----------------|----------------|-------------------|---------------------|------------------|
| Application | | Application group | | |
| tcp | 10.1.1.2 | 81 | 20.1.1.2 | 32813 |
| junos:ftp [63] | | unknown [1023] | | |
| tcp | 20.1.1.2 | 32813 | 10.1.1.2 | 81 |
| junos:ftp [63] | | unknown [1023] | | |

show services local-policy-decision-function statistics

| | |
|---------------------------------|--|
| Syntax | <code>show services local-policy-decision-function statistics (interface <i>interface-name</i> subscriber <i>subscriber-name</i>)</code> |
| Release Information | Command introduced in Junos OS Release 9.5. |
| Description | Display local-policy-decision-function (L-PDF) statistics. |
| Options | <p><code>interface <i>interface-name</i></code>—Display L-PDF statistics for the specified interface(s) only.</p> <p><code>subscriber <i>subscriber-name</i></code>—Display L-PDF statistics for the specified subscriber(s) only.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show services local-policy-decision-function statistics by interface on page 2116</p> <p>show services local-policy-decision-function statistics by subscriber on page 2117</p> |
| Output Fields | <p>Table 117 on page 2116 lists the output fields for the show services local-policy-decision-function statistics command. Output fields are listed in the approximate order in which they appear.</p> |

Table 117: show services local-policy-decision-function statistics Output Fields

| Field Name | Field Description |
|-----------------------|-------------------------------|
| Interface | Interface name. |
| service-set | Service set name. |
| service-set-interface | Service set interface name. |
| Application group | Application group identifier. |
| Application | Application name. |
| Packets in | Number of ingress packets. |
| Bytes in | Number of ingress bytes. |
| Packets out | Number of egress packets. |
| Bytes out | Number of egress bytes. |

Sample Output

show services local-policy-decision-function statistics by interface

```
user@host> show services local-policy-decision-function statistics interface ge-1/1/0
```

Interface: ge-1/1/0.0

service-set: IDP

service-set interface: ms-2/0/0

| Application group | Application | Packets in | Bytes in |
|-------------------|----------------|------------|----------|
| Packets out | Bytes out | | |
| | junos:ftp [63] | 5 | 334 |
| 6 | 346 | | |

show services local-policy- decision-function statistics by subscriber

user@host> show services local-policy-decision-function statistics subscriber user@juniper.net

Service-set-interface: ms-1/3/0

Service set: aacl-svc-set

Application-aware-access-list statistics

| Application group | Packets in | Bytes in | Packets out | Bytes |
|-------------------|------------|----------|-------------|-------|
| P2P | | 400 | 32025 | 200 |
| | 16284 | | | |
| FTP | | 20000 | 5231000 | 100 |
| | 8700 | | | |

Link and Multilink Services Operational Commands

- [show interfaces \(Link Services\)](#)
- [show interfaces \(Link Services IQ\)](#)
- [show interfaces \(Multilink Services\)](#)

show interfaces (Link Services)

Syntax For Multilink Frame Relay user-to-user network-to-network interface (UNI NNI):

```
show interfaces interface-type :channel
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

For Multilink Frame Relay end-to-end:

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Release Information Command introduced before Junos OS Release 7.4.

Description Display status information about the specified link services interface.

Options *interface-type*—On M Series and T Series routers, the interface type is *ls-fpc/pic/port*.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level view

List of Sample Output [show interfaces extensive \(MFR UNI NNI\) on page 2126](#)
[show interfaces extensive \(MFR End-to-End\) on page 2128](#)

Output Fields [Table 118 on page 2118](#) lists the output fields for the **show interfaces** (link services) command. Output fields are listed in the approximate order in which they appear.

Table 118: Link Services show interfaces Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|---------------------------------|-----------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |

Table 118: Link Services show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|--|------------------------------|
| Enabled | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface's index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| Link-level type | Encapsulation being used on the physical interface:
Multilink-Frame-Relay-UNI-NNI (default), LinkService , Frame-relay , Frame-relay-ccc , or Frame-relay-tcc . | All levels |
| MTU | Maximum transmission unit size on the physical interface. | All levels |
| Device flags | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) . | detail extensive none |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Link flags | Information about the link. Possible values are described in the “Link Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Hold-times | Current interface hold time up and hold time down, in milliseconds, in the format Up <i>n</i> ms, Down <i>n</i> ms . | detail extensive |

Table 118: Link Services show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|--|-----------------------|
| Multilink Frame Relay UNI NNI bundle options | <p>Multilink Frame Relay UNI NNI only) Configured information about Multilink Frame Relay bundle options.</p> <ul style="list-style-type: none"> • Device type—DCE (Data Communication Equipment) or DTE (Data Terminal Equipment). • MRRU—Configured size of the maximum received reconstructed unit (MRRU): 1500 to 4500 bytes. The default is 1524 bytes. • Fragmentation threshold—Configured fragmentation threshold: 128 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation. • Red differential delay limit—Red differential delay limit among bundle links has been reached, indicating an action will occur. • Yellow differential delay limit—Yellow differential delay among bundle links has been reached, indicating a warning will occur. • Red differential delay action—Type of actions taken when the red differential delay exceeds the red limit: Disable link transmit or Remove link from service. • Reassembly drop timer—Drop timeout value to provide a recovery mechanism if individual links in the link services bundle drop one or more packets: 1 through 127 milliseconds. By default, the drop timeout parameter is 0 (disabled). A value that is under 5 ms is not recommended. • Links needed to sustain bundle—Minimum number of links to sustain the bundle: 1 through 8. • LIP Hello timer—Link Interleaving Protocol hello timer: 1 through 180 seconds. <ul style="list-style-type: none"> • Acknowledgement timer—Maximum period to wait for an add link acknowledgement, hello acknowledgement, or remove link acknowledgement: 1 through 10 seconds. • Acknowledgement retries—Number of retransmission attempts to be made for consecutive hello or remove link messages after the expiration of the acknowledgement timer: 1 through 5. | detail extensive none |
| Multilink Frame Relay UNI NNI bundle options (continued) | <ul style="list-style-type: none"> • Bundle class—Bundle class ID. • LMI type—Multilink Frame Relay UNI NNI LMI type: ANSI or Q.933 ANNEX A. <ul style="list-style-type: none"> • T391 LIV polling timer—Multilink Frame Relay UNI NNI Full status polling counter: 1 through 255, with a default value of 6. • T392 polling verification timer—Multilink Frame Relay UNI NNI LMI error threshold. The number of errors required to bring down the link, within the event count specified by N393. The range is 1 through 10, with a default value of 3. • N391 full status polling count—Multilink Frame Relay UNI NNI Full status polling counter: 1 through 255. • N392 error threshold—Multilink Frame Relay UNI NNI LMI error threshold: 1 through 10. • N393 monitored event count—Multilink Frame Relay UNI NNI LMI monitored event count: 1 through 10, with a default value of 4. | detail extensive none |
| Traffic statistics | <p>Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router.</p> | detail extensive |

Table 118: Link Services show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|---|-----------------------|
| Multilink Frame Relay UNI NNI bundle errors | Information about Multilink Frame Relay bundle errors. <ul style="list-style-type: none"> • Packet drops—Number of packets dropped. • Fragment drops—Number of fragments dropped. • MRRU exceeded—Number of times a packet was dropped because the configured MRRU value was exceeded. • Exception events—Exception events counter. | detail extensive |
| Multilink Frame Relay UNI NNI bundle statistics | Information about Multilink Frame Relay bundles. <ul style="list-style-type: none"> • Fragments—Bundle fragment information. <ul style="list-style-type: none"> • Input—Total number and rate of frames and packets received, in Frames, fps (frames per second), Bytes, and bps (bits per second). • Output—Total number and rate of frames and packets transmitted, in Frames, fps, Bytes, and bps. • Packets—Bundle packet information. <ul style="list-style-type: none"> • Input—Total number and rate of frames and packets received, in Frames, fps (frames per second), Bytes, and bps (bits per second). • Output—Total number and rate of frames and packets transmitted, in Frames, fps, Bytes, and bps. | detail extensive |
| Multilink Frame Relay UNI NNI bundle links information | <ul style="list-style-type: none"> • Active bundle links—Number of bundle links that are currently active. • Removed bundle links—Number of bundle links that have been removed (RED differential delay action). • Disabled bundle links—Number of bundle links that have been disabled (RED differential delay action). | detail extensive none |
| Multilink Frame Relay UNI NNI active bundle links statistics | (Multilink Frame Relay UNI NNI only) Display information for each active bundle link. <ul style="list-style-type: none"> • Frames—Number of multilink control frames received on this bundle link. • fps—Rate of multilink control frames received on this bundle link (in frames per second). • Bytes—Number of bytes received on this bundle link. • bps—Number of bits per second received on this bundle link. • interface-name—Name of the bundle link interface. • Input—Total number and rate of frames and packets received. • Output—Total number and rate of frames and packets transmitted. • Current differential delay—Compare this bundle link's round trip time to the average of all bundle links' round trip times in ms (milliseconds). • Recent high differential delay—Highest differential delay value from the latest 10 intervals, in milliseconds. • Times over red diff delay—Number of times this bundle link exceeded the configured red differential delay limit. • Times over yellow diff delay—Number of times this bundle link exceeded the configured yellow differential delay limit. | detail extensive |

Table 118: Link Services show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|---|------------------|
| Multilink Frame Relay UNI NNI active bundle links statistics
(continued) | <ul style="list-style-type: none"> LIP—Link Interleaving Protocol information. Rcv—Number of messages received. Xmt—Number of messages transmitted. add_lnk—ADD_LINK message notifies the peer endpoint that the local endpoint supports frame processing. It is generated on both ends of a bundle link when a bundle link endpoint is ready to become operational. lnk_ack—ADD_LINK_ACK message notifies the peer that the local router has received a valid ADD_LINK message. lnk_rej—ADD_LINK_REJ message notifies the peer that the local router has received an invalid ADD_LINK message. hello—HELLO message notifies the peer that the local router is up. Both ends of a link bundle generate this message. hel_ack—HELLO_ACK message notifies the peer that the local router has received a valid HELLO message. lnk_rem—REMOVE_LINK message notifies the peer that the local router has received a REMOVE_LINK message. rem_ack—REMOVE_LINK_ACK message notifies the peer that the local router has received a valid ADD_LINK message. | detail extensive |
| Frame exceptions | <p>For Multilink Frame Relay end-to-end only. Information about framing exceptions. Includes events recorded under Exception Events for each logical interface.</p> <ul style="list-style-type: none"> Oversized frames—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb (kilobits). Errored input frames—Number of input frame errors. Input on disabled link/bundle—Number of frames received on disabled links. These frames can result either from an inconsistent configuration, or from a bundle or link being brought up or down with traffic actively flowing through it. Output for disabled link/bundle—Number of frames sent for a disabled or unavailable link. These frames can result either from an inconsistent configuration, or from a bundle being brought up or down while traffic is flowing through it. Queuing drops—Total number of packets dropped before traffic enters the link services IQ interface. Indicates that the interface is becoming oversubscribed. | detail extensive |
| Buffering exceptions | <p>For Multilink Frame Relay end-to-end only. Information about buffering exceptions. Includes events recorded under Exception Events for each logical interface:</p> <ul style="list-style-type: none"> Packet data buffer overflow—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services interface capacity. Fragment data buffer overflow—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services interface capacity. Check the logical interface exception event counters to determine which bundle is responsible. | detail extensive |

Table 118: Link Services show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|---|-------------------------|
| Assembly exceptions | <p>For Multilink Frame Relay end-to-end only. Information about assembly exceptions. Includes events recorded under Exception Events for each logical interface.</p> <p>An assembly exception does not necessarily indicate an operational problem with the physical link services interface itself. If multilink-encapsulated traffic is dropped or reordered after a sequence number has been assigned, the assembling multilink interface records one or more exception events. The multilink interface can drop multilink-encapsulated fragments itself as a result. Any multilink packets or fragments dropped by the physical link services interface itself result in packet or fragment drop counts on individual logical interfaces. If the logical interface drop counts are zero, but exception events are seen, the most likely cause is a problem with the individual link interfaces. Even if the logical interface fragment drop counts are nonzero, excess differential delay or traffic losses on individual interfaces can be the root cause.</p> | detail extensive |
| Assembly exceptions (continued) | <ul style="list-style-type: none"> • Fragment timeout—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. These events do not necessarily indicate any problem with the operation of the physical link services interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • Missing sequence number—A gap was detected in the sequence numbers of fragments on a bundle. These events do not necessarily indicate any problem with the operation of the physical link services interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • Out-of-order sequence number—Two frames with out-of-order sequence numbers occurred within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. Check the logical interface exception event counters to determine which bundle is responsible. • Out-of-range sequence number—Frame was received with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. Check the logical interface exception event counters to determine which bundle is responsible. | detail extensive |
| Hardware errors | <p>For Multilink Frame Relay end-to-end only. Information about hardware errors:</p> <ul style="list-style-type: none"> • Data memory error—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. • Control memory error—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. | detail extensive |
| Logical Interface | | |
| Logical interface | Name of the logical interface. | All levels |

Table 118: Link Services show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|---|------------------------------|
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| Flags | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Encapsulation | Encapsulation being used: PPP, Multilink - FR or Multilink - PPP | All levels |
| Bandwidth | Speed at which the interface is running. | All levels |
| Bundle options | For Multilink Frame Relay end-to-end interfaces only: <ul style="list-style-type: none"> • MRRU—Configured size of the maximum received reconstructed unit (MRRU): 1500 to 4500 bytes. The default is 1524 bytes. • Drop timer period—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: 1 through 127 milliseconds. Values under 5 milliseconds are not recommended. The default setting is 0, which disables the timer. • Sequence number format—(MLPPP) Short sequence number header format. • Fragmentation threshold—Configured fragmentation threshold: 128 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation. • Links needed to sustain bundle—Minimum number of links to sustain the bundle: 1 through 8. • Interleave fragments—State of the process that interleaves long packets with high-priority ones. Only Disabled is currently supported. • Remote MRRU—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed. | detail extensive none |
| Bundle status (MLPPP) or Multilink class status (MC-MLPPP) | Information about bundle status: <ul style="list-style-type: none"> • Remote MRRU—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed. • Received sequence number—Sequence number for received packets. • Transmit sequence number—Sequence number for transmitted packets. • Packet drops—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • Fragment drops—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. | detail extensive none |

Table 118: Link Services show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|---|------------------------------|
| Bundle status (MLPPP) or Multilink class status (MC-MLPPP) (continued) | <ul style="list-style-type: none"> • MRRU exceeded—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • Fragment timeout—Drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • Missing sequence number—Gap detected in the sequence numbers of fragments on a bundle. • Out-of-order sequence number—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • Out-of-range sequence number—Frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • Packet data buffer overflow—Packet buffer memory full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • Fragment data buffer overflow—Fragment buffer memory full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. | detail extensive none |
| Bundle errors | <p>Information about bundle errors.</p> <ul style="list-style-type: none"> • Packet drops—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. • Fragment drops—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. • MRRU exceeded—Number of reassembled packets exceeding the MRRU. • Exception events—Number of exceptional events encountered other than MRRU exceeded errors. These events are categorized under the physical interface: Frame exceptions, Buffering exceptions, and Fragment exceptions. Exception events do not necessarily indicate that the multilink interface is not operating properly. Individual link failures can produce exceptional events. | detail extensive |
| Statistics | <p>Information about fragments and packets received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router.</p> <ul style="list-style-type: none"> • Bundle—Information about bundles. • Link—Information about links used in the multilink operation. | detail extensive |
| Protocol | Protocol family configured on the logical interface. | detail extensive none |
| MTU | MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked Adjusted . | detail extensive none |

Table 118: Link Services show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|--|-----------------------|
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route table | Routing table in which this address exists. For example, Route table:0 refers to inet.0. | detail extensive |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Addresses, Flags | Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Destination | IP address of the remote side of the connection. | detail extensive none |
| Local | IP address of the logical interface. | detail extensive none |
| Broadcast | Broadcast address on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |

Sample Output

show interfaces extensive (MFR UNI NNI)

```

user@host> show interfaces ls-1/3/0:0 extensive
Physical interface: ls-1/3/0:0, Enabled, Physical link is Up
Interface index: 25, SNMP ifIndex: 35, Generation: 124
Link-level type: Multilink-FR-UNI-NNI, MTU: 1524
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps
Last flapped   : 2002-11-01 15:26:25 PST (00:34:49 ago)
Statistics last cleared: Never
Link flags     : None
Hold-times     : Up 0 ms, Down 0 ms
Multilink Frame Relay UNI NNI bundle options:
  Device type           DTE
  MRRU                  1524
  Fragmentation threshold 1500
  Red differential delay limit 10
  Yellow differential delay limit 6
  Red differential delay action Disable link transmit
  Reassembly drop timer 0
  Links needed to sustain bundle 1
  LIP Hello timer       10
    Acknowledgement timer 4
    Acknowledgement retries 2
  Bundle class          A
  LMI type               Q.933 Annex A
    T391 LIV polling timer 10
    T392 polling verification timer 15
    N391 full status polling count 6
    N392 error threshold 3
    N393 monitored event count 4

```



```

Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Multilink Frame Relay UNI NNI bundle errors:
Packet drops 0 (0 bytes)
Fragment drops 0 (0 bytes)
MRRU exceeded 0
Exception events 0
Multilink Frame Relay UNI NNI bundle statistics
      Frames      fps      Bytes      bps
Fragments:
Input : 0 0 0 0
Output: 824 0 17304 320
Packets:
Input : 0 0 0 0
Output: 824 0 17304 320
Multilink Frame Relay UNI NNI bundle links info:
Active bundle links 4
Removed bundle links 0
Disabled bundle links 0
Multilink Frame Relay UNI NNI active bundle links statistics:
      Frames      fps      Bytes      bps
t1-0/2/0:0.0
Input : 0 0 0 0
Output: 206 0 4326 80
Current differential delay 0.2 ms
Recent high differential delay 3.8 ms
Times over red diff delay 0
Times over yellow diff delay 0
LIP:add_lnk lnk_ack lnk_rej hello hel_ack lnk_rem rem_ack
Rcv: 2 2 0 206 207 0 0
Xmt: 2 1 0 207 206 0 0
t1-0/2/0:1.0
Input : 0 0 0 0
Output: 206 0 4326 80
Current differential delay 0.2 ms
Recent high differential delay 3.7 ms
Times over red diff delay 0
Times over yellow diff delay 0
LIP:add_lnk lnk_ack lnk_rej hello hel_ack lnk_rem rem_ack
Rcv: 2 2 0 206 207 0 0
Xmt: 2 1 0 207 206 0 0
t1-0/2/0:2.0
Input : 0 0 0 0
Output: 206 0 4326 80
Current differential delay 0.4 ms
Recent high differential delay 3.8 ms
Times over red diff delay 0
Times over yellow diff delay 0
LIP:add_lnk lnk_ack lnk_rej hello hel_ack lnk_rem rem_ack
Rcv: 2 2 0 206 207 0 0
Xmt: 2 1 0 207 206 0 0
t1-0/2/0:3.0
Input : 0 0 0 0
Output: 206 0 4326 80
Current differential delay 0.3 ms
Recent high differential delay 3.8 ms
Times over red diff delay 0

```

```

Times over yellow diff delay      0
LIP:add_lnk lnk_ack lnk_rej      hello hel_ack lnk_rem rem_ack
Rcv:      2      2      0      206      207      0      0
Xmt:      2      1      0      207      206      0      0
Logical interface ls-1/3/0:0.0 (Index 5) (SNMP ifIndex 28) (Generation 10)
Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI
Bandwidth: 622080kbps
Bundle errors:
  Packet drops                    0 (0 bytes)
  Fragment drops                  0 (0 bytes)
  MRRU exceeded                   0
  Exception events                0
Statistics      Frames      fps      Bytes      bps
Bundle:
  Fragments:
    Input :      0      0      0      0
    Output:     824      0     17304     320
  Packets:
    Input :      0      0      0      0
    Output:     824      0     17304     320
Link:
  t1-0/2/0:0.0
    Input :      0      0      0      0
    Output:     206      0     4326      80
  t1-0/2/0:1.0
    Input :      0      0      0      0
    Output:     206      0     4326      80
  t1-0/2/0:2.0
    Input :      0      0      0      0
    Output:     206      0     4326      80
  t1-0/2/0:3.0
    Input :      0      0      0      0
    Output:     206      0     4326      80
Protocol inet, MTU: 1500 [Adjusted], Generation: 15, Route table: 0
Flags: User-MTU, MTU-Protocol-Adjusted
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 1.1.1.1, Local: 1.1.1.2, Broadcast: Unspecified,
  Generation: 10

```

show interfaces extensive (MFR End-to-End)

```

user@host> show interfaces ls-0/3/0 extensive
Physical interface: ls-0/3/0, Enabled, Physical link is Up
Interface index: 264, SNMP ifIndex: 104, Generation: 525
Link-level type: LinkService, MTU: 1524
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps
Last flapped   : 2002-10-16 17:53:49 PDT (00:22:00 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :      73471      264 bps
Output bytes :     80335      0 bps
Input packets:      822      0 pps
Output packets:     819      0 pps
Frame exceptions:
Oversized frames      0
Errored input frames  0
Input on disabled link/bundle 0
Output for disabled link/bundle 4
Queuing drops         3
Buffering exceptions:

```

```

    Packet data buffer overflow      0
    Fragment data buffer overflow    0
Assembly exceptions:
    Fragment timeout                 0
    Missing sequence number          0
    Out-of-order sequence number     0
    Out-of-range sequence number     0
Hardware errors (sticky):
    Data memory error                0
    Control memory error              0
Logical interface ls-0/3/0.0 (Index 5) (SNMP ifIndex 527) (Generation 47)
Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
Bandwidth: 1536kbps
Bundle options:
    MRRU                            1524
    Drop timer period                 0
    Sequence number format            long (24 bits)
    Fragmentation threshold           0
    Links needed to sustain bundle    1
    Interleave fragments              Disabled
Bundle status:
    Remote MRRU                      1500
    Received sequence number          0x19ec14
    Transmit sequence number          0x38cfa8
    Packet drops                      0 (0 bytes)
    Fragment drops                    0 (0 bytes)
    MRRU exceeded                     0
    Fragment timeout                   0
    Missing sequence number            0
    Out-of-order sequence number       0
    Out-of-range sequence number       0
    Packet data buffer overflow        0
    Fragment data buffer overflow      0
Bundle errors:
    Packet drops                      2 (68 bytes)
    Fragment drops                    0 (0 bytes)
    MRRU exceeded                     0
    Exception events                   0
Statistics      Frames      fps      Bytes      bps
Bundle:
  Fragments:
    Input :      172         0      15544      288
    Output:      165         0      16645       0
  Packets:
    Input :      143         0      12885      288
    Output:      134         0      12276       0
Link:
  t1-0/0/0.0
    Input :      143         0      12885      288
    Output:      134         0      12276       0
Protocol inet, MTU: 1500, Generation: 76, Route table: 0
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.16.1.2, Local: 10.16.1.1, Broadcast:
  Unspecified, Generation: 81
Protocol iso, MTU: 1500 [Adjusted], Generation: 77, Route table: 0
  Flags: Is-Primary
Protocol inet6, MTU: 1500, Generation: 78, Route table: 0
  Flags: Is-Primary
  Addresses, Flags: Is-Default Is-Preferred Is-Primary
  Destination: 8016::1:0/126, Local: 8016::1:1,

```

Broadcast: Unspecified, Generation: 83
Addresses, Flags: Is-Preferred
Destination: fe80::/64, Local: fe80::2a0:a5ff:fe12:4777,
Broadcast: Unspecified,
Generation: 85

show interfaces (Link Services IQ)

| | |
|---------------------------------|--|
| Syntax | <pre>show interfaces lsq-<i>fpc/pic/port</i> <brief detail extensive terse> <descriptions> <l2-statistics> <media> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>l2-statistics option introduced with Junos OS Release 12.1.</p> |
| Description | (M Series, MX Series, and T Series routers only) Display status information about the specified link services intelligent queuing (IQ) interface. |
| Options | <p>lsq-<i>fpc/pic/port</i>—Display standard status information about the specified link services IQ interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>l2-statistics—(Optional) Display Layer 2 queue statistics for Multilink Point-to-Point Protocol (MLPPP), FRF.15, and FRF.16 bundles.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Additional Information | Link services IQ interfaces are similar to link services interfaces. The important difference is that link services IQ interfaces fully support Junos OS class-of-service (CoS) components. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • Link and Multilink Services Overview on page 711 • Multilink Interfaces on Channelized MICs Overview on page 715 |
| List of Sample Output | <p>show interfaces extensive (MLPPP on Link Services IQ) on page 2146</p> <p>show interfaces extensive (Multiclass MLPPP on Link Services IQ) on page 2147</p> <p>show interfaces extensive (MLPPP on Link Services IQ Bundle) on page 2149</p> <p>show interfaces extensive (MFR on Link Services IQ Bundle) on page 2150</p> <p>show interfaces (Multiclass MLPPP on Link Services IQ) on page 2152</p> |

Output Fields Table 63 on page 1886 lists the output fields for the **show interfaces** (link services IQ) command. Output fields are listed in the approximate order in which they appear.

Table 119: show interfaces (Link Services IQ) Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|--|------------------------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| Link-level type | Encapsulation being used on the physical interface:
Multilink-Frame-Relay-UNI-NNI Multilink-Frame-Relay-UNI-NNI (default),
LinkService , Frame-relay , Frame-relay-ccc , or Frame-relay-tcc . | All levels |
| MTU | Maximum transmission unit size on the physical interface. | All levels |
| Device flags | Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) . | detail extensive none |

Table 119: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|--|------------------------------|
| Multilink Frame Relay UNI NNI bundle options | <p>(Multilink Frame Relay UNI NNI only) Configured information about Multilink Frame Relay bundle options.</p> <ul style="list-style-type: none"> • Device type—DCE (data communication equipment) or DTE (data terminal equipment). • MRRU—Configured size of the maximum received reconstructed unit (MRRU): 1500 to 4500 bytes. The default is 1524 bytes. • Bandwidth—Speed at which the interface is running. • Fragmentation threshold—Configured fragmentation threshold: 128 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation. • Red differential delay limit—Red differential delay limit among bundle links has been reached, indicating an action will occur. • Yellow differential delay limit—Yellow differential delay among bundle links has been reached, indicating a warning will occur. • Red differential delay action—Type of actions taken when the red differential delay exceeds the red limit: <i>Disable link transmit</i> or <i>Remove link from service</i>. • Link layer overhead—Percentage of bundle bandwidth to be set aside for link layer overhead. • Reassembly drop timer—Drop timeout value to provide a recovery mechanism if individual links in the link services bundle drop one or more packets: 1 through 127 milliseconds. By default, the drop timeout parameter is 0 (disabled). A value under 5 ms is not recommended. • Links needed to sustain bundle—Minimum number of links to sustain the bundle: 1 through 8. • LIP Hello timer—Link Interleaving Protocol hello timer: 1 through 180 seconds. <ul style="list-style-type: none"> • Acknowledgement timer—Maximum period to wait for an add link acknowledgement, hello acknowledgement, or remove link acknowledgement: 1 through 10 seconds. • Acknowledgement retries—Number of retransmission attempts to be made for consecutive hello or remove link messages after the expiration of the acknowledgement timer: 1 through 5. | detail extensive none |

Table 119: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|--|-----------------------|
| Multilink Frame Relay UNI NNI bundle options (continued) | <ul style="list-style-type: none"> • Bundle class—Bundle class ID. • LMI type—Multilink Frame Relay UNI NNI LMI type: ANSI, Q.933 ANNEX A, or Consortium. <ul style="list-style-type: none"> • T391 LIV polling timer—Multilink Frame Relay UNI NNI Full status polling counter: 1 through 255, with a default value of 6. • T392 polling verification timer—Multilink Frame Relay UNI NNI LMI error threshold. The number of errors required to bring down the link, within the event count specified by <i>N393</i>. The range is 1 through 10, with a default value of 3. • N391 full status polling count—Multilink Frame Relay UNI NNI Full status polling counter: 1 through 255. • N392 error threshold—Multilink Frame Relay UNI NNI LMI error threshold: 1 through 10. • N393 monitored event count—Multilink Frame Relay UNI NNI LMI monitored event count: 1 through 10, with a default value of 4. • Consortium LMI Settings <ul style="list-style-type: none"> • n391dte—DTE full status polling interval in seconds: 1 through 255. • n392dce—DCE error threshold: 1 through 10. • n392dte—DTE error threshold: 1 through 10. • n393dce—DCE monitored event count: 1 through 10. • n393dte—DTE monitored event count: 1 through 10. • t391dte—DTE polling verification timer (in seconds): 5 through 30. • t392dce—DCE polling verification timer (in seconds): 5 through 30. | detail extensive none |
| LMI | <p>Local Management Interface packet statistics:</p> <ul style="list-style-type: none"> • Input—Number of packets arriving on the interface (nn) and timestamp of the most recent packet arrival, in the format:
Input: nn (last seen hh:mm:ss ago) • Output—Number of packets sent out on the interface (nn) and how much time has passed since the last packet was sent, in the format:
Output: nn (last seen hh:mm:ss ago) | detail extensive none |
| DTE Statistics | <p>Statistics about information transferred from the data terminal equipment (DTE) to the data communications equipment (DCE).</p> <ul style="list-style-type: none"> • Enquiries sent—Number of link status enquiries sent from the DTE to the DCE. • Full enquiries sent—Number of full enquiries sent from the DTE to the DCE. • Enquiry responses received—Number of enquiry responses received by the DCE from the DTE. • Full enquiry responses received—Number of full enquiry responses received by DCE from the DTE. | detail extensive none |

Table 119: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|--|------------------------------|
| DCE Statistics | <p>Statistics about information transferred from the DCE to the DTE.</p> <ul style="list-style-type: none"> • Enquiries received—Number of enquiries received by the DCE from the DTE. • Full enquiries received—Number of full enquiries received by the DCE from the DTE. • Enquiry responses sent—Number of enquiry responses sent from the DCE to the DTE. • Full enquiry responses sent—Number of full enquiry responses sent from the DCE to the DTE. | detail extensive none |
| Common Statistics | <p>Statistics about messages sent between the DTE and the DCE.</p> <ul style="list-style-type: none"> • Unknown messages received—Number of received packets that do not fall into any other category. • Asynchronous updates received—Number of link status peer changes received. • Out-of-sequence packets received—Number of packets for which the sequence of the packets received is different from the expected sequence. • Keepalive responses timed out—Number of keepalive responses that time out when no Local Management Interface (LMI) packet was reported for n392dte or n393dce intervals. (See <i>LMI settings</i>.) | |
| Traffic statistics | <p>Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the Packet Forwarding Engine (PFE). Input traffic refers to the fragments received by the ingress PFE, which get assembled into Layer 3 input packets. Output packets refer to the IP packets transmitted out of the ingress PFE to the LSQ, which get segmented into output fragments.</p> | detail extensive |
| DLCInn | <p>Data-link connection identifier (DLCI) number of the logical interface. The following information is displayed.</p> <ul style="list-style-type: none"> • Flags—Values are: <ul style="list-style-type: none"> • Active—Set when the link is active and the DTE and DCE are exchanging information. • Down—Set when the link is active, but no information is received from the DTE. • DCE unconfigured—Set when the corresponding DLCI in the DCE is not configured. • Configured—Set when the corresponding DLCCI is configured. • DCE-Configured—Displayed when the command is issued from the DTE. | |
| DLCI Statistics | <p>(Frame Relay) Data-link connection identifier (DLCI) statistics.</p> <ul style="list-style-type: none"> • Active DLCI—Number of active DLCIs. • Inactive DLCI—Number of inactive DLCIs. | |
| Input rate | (Redundant LSQ) Rate of bits and packets received on the interface. | None specified |
| Output rate | (Redundant LSQ) Rate of bits and packets transmitted on the interface. | None specified |

Table 119: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|-------------------------|
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Traffic statistics | Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. | detail extensive |
| Frame exceptions | <p>Information about framing exceptions. Includes events recorded under Exception Events for each logical interface.</p> <ul style="list-style-type: none"> • Oversized frames—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb (kilobits). • Errored input frames—Number of input frame errors. • Input on disabled link/bundle—Number of frames received on disabled links. These frames can result either from an inconsistent configuration, or from a bundle or link being brought up or down with traffic actively flowing through it. • Output for disabled link/bundle—Number of frames sent for a disabled or unavailable link. These frames can result either from an inconsistent configuration, or from a bundle being brought up or down while traffic is flowing through it. • Queuing drops—Total number of packets dropped before traffic enters the link services IQ interface. Indicates that the interface is becoming oversubscribed. | extensive |
| Buffering exceptions | <p>Information about buffering exceptions. Includes events recorded under Exception Events for each logical interface:</p> <ul style="list-style-type: none"> • Packet data buffer overflow—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • Fragment data buffer overflow—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. Check the logical interface exception event counters to determine which bundle is responsible. | extensive |

Table 119: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|---|------------------------------|
| Assembly exceptions | <p>(Multilink Frame Relay end-to-end only) Information about assembly exceptions. Includes events recorded under Exception Events for each logical interface.</p> <p>An assembly exception does not necessarily indicate an operational problem with the physical link services IQ interface itself. If multilink-encapsulated traffic is dropped or reordered after a sequence number has been assigned, the interface records one or more exception events. The physical interface can drop multilink-encapsulated fragments itself as a result. Any multilink packets or fragments dropped by the interface itself result in packet or fragment drop counts on individual logical interfaces. If the logical interface drop counts are zero, but exception events are seen, the most likely cause is a problem with the individual link interfaces. Even if the logical interface fragment drop counts are nonzero, excess differential delay or traffic losses on individual interfaces can be the root cause.</p> <ul style="list-style-type: none"> • Fragment timeout—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • Missing sequence number—A gap was detected in the sequence numbers of fragments on a bundle. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • Out-of-order sequence number—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the link services IQ interface unable to correctly process the resulting stream. Check the logical interface exception event counters to determine which bundle is responsible. • Out-of-range sequence number—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. Check the logical interface exception event counters to determine which bundle is responsible. | extensive |
| Hardware errors (sticky) | <p>(Multilink Frame Relay end-to-end only) Information about hardware errors:</p> <ul style="list-style-type: none"> • Data memory error—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. • Control memory error—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. | extensive |
| Egress queues | Total number of egress queues supported on the specified interface. | detail extensive none |

Table 119: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------|---|------------------------------|
| Queue counters | Queue number and its associated user-configured forwarding class name. <ul style="list-style-type: none"> Queued packets—Number of queued packets. Transmitted packets—Number of transmitted packets. Dropped packets—Number of packets dropped by the ASIC's RED mechanism. | detail extensive none |
| Logical Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| Flags | Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Encapsulation | Encapsulation being used: PPP or Multilink PPP. | All levels |
| Bandwidth | Speed at which the interface is running. | All levels |
| Bundle options | (Multilink Frame Relay end-to-end interfaces only) <ul style="list-style-type: none"> MRRU—Configured size of the maximum received reconstructed unit (MRRU): 1500 through 4500 bytes. The default is 1504 bytes. Drop timer period—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: 0 through 2000 milliseconds. Values under 5 ms are not recommended. The default setting is 0, which disables the timer. Sequence number format—Short sequence number header format (MLPPP only). Fragmentation threshold—Configured fragmentation threshold: 64 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation. Links needed to sustain bundle—Minimum number of links to sustain the bundle: 1 through 8. Multilink classes—Number of multilink classes negotiated. Link layer overhead—Percentage of bundle bandwidth to be set aside for link-layer overhead. | detail extensive none |

Table 119: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|---|------------------------------|
| Bundle status
(MLPPP) or
Multilink class status (Multiclass MLPPP) | Information about bundle status: <ul style="list-style-type: none"> • Remote MRRU—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed. • Received sequence number—Sequence number for received packets. • Transmitted sequence number—Sequence number for transmitted packets. • Packet drops—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • Fragment drops—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully, but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • MRRU exceeded—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • Fragment timeout—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • Missing sequence number—A gap was detected in the sequence numbers of fragments on a bundle. • Out-of-order sequence number—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • Out-of-range sequence number—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • Packet data buffer overflow—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • Fragment data buffer overflow—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. | detail extensive none |

Table 119: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|---|------------------------------|
| Statistics | <p>Information about fragments and packets received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).</p> <ul style="list-style-type: none"> • Bundle—Information for each active bundle link. <ul style="list-style-type: none"> • Fragments: Input and Output—Total number and rate of fragments received and transmitted. • Packets: Input and Output—Total number and rate of packets received and transmitted. • Multilink class—(Multiclass MLPPP only) Information about multiclass links used in the multilink operation. • Link—Information about links used in the multilink operation. <ul style="list-style-type: none"> • Link name—Interface name of the link services IQ channel and state information (physical link up or down). • Input and Output—Total number and rate of fragments and packets received and transmitted. | detail extensive |
| NCP state | <p>(PPP) Network Control Protocol state.</p> <ul style="list-style-type: none"> • Conf-ack-received—Acknowledgement was received. • Conf-ack-sent—Acknowledgement was sent. • Conf-req-sent—Request was sent. • Down—NCP negotiation is incomplete (not yet completed or has failed). • Not-configured—NCP is not configured on the interface. • Opened—NCP negotiation is successful. | detail extensive none |
| Protocol | Protocol family configured on the logical interface. | detail extensive none |
| MTU | MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked Adjusted . | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route Table | Routing table in which this address exists. For example, Route table:0 refers to inet.0. | detail extensive |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Addresses, Flags | Information about the addresses configured on the logical interface. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Destination | IP address of the remote side of the connection. | detail extensive none |

Table 119: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|---|------------------------------|
| Local | IP address of the logical interface. | detail extensive none |
| Broadcast | Broadcast address on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| MLPPP Bundle Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| Flags | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| SNMP-Traps | SNMP trap notifications are enabled. | All levels |
| Encapsulation | Encapsulation being used: PPP, Multilink PPP, or Multilink-FR. | All levels |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) . | detail extensive none |
| Bandwidth | Speed at which the interface is running. | All levels |
| Bundle links information | Information about the bundled links. <ul style="list-style-type: none"> • Active bundle links—Number of active links. • Removed bundle links—Information about links used in the multilink operation. • Disabled bundle links—Number of disabled links. | detail extensive none |

Table 119: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------|--|------------------------------|
| Bundle options | <p>(Multilink Frame Relay end-to-end interfaces only)</p> <ul style="list-style-type: none"> • MRRU—Configured size of the maximum received reconstructed unit (MRRU): 1500 through 4500 bytes. The default is 1504 bytes. • Drop timer period—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: 0 through 2000 milliseconds. Values under 5 ms are not recommended. The default setting is 0, which disables the timer. • Inner PPP Protocol field compression—Inner PPP protocol compression is enabled or disabled. • Sequence number format—Short sequence number header format (MLPPP only). • Fragmentation threshold—Configured fragmentation threshold: 64 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation. • Links needed to sustain bundle—Minimum number of links to sustain the bundle: 1 through 8. • Multilink classes—Number of multilink classes negotiated. • Link layer overhead—Percentage of bundle bandwidth to be set aside for link-layer overhead. | detail extensive none |

Table 119: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|--|------------------------------|
| Bundle status
(MLPPP) | <p>Information about bundle status:</p> <ul style="list-style-type: none"> • Received sequence number—Sequence number for received packets. • Transmit sequence number—Sequence number for transmitted packets. • Packet drops—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • Fragment drops—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • MRRU exceeded—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • Fragment timeout—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • Missing sequence number—A gap was detected in the sequence numbers of fragments on a bundle. • Out-of-order sequence number—Two frames with out-of-order sequence numbers occurred within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • Out-of-range sequence number—A frame was received with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • Packet data buffer overflow—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • Fragment data buffer overflow—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. | detail extensive none |

Table 119: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------|--|------------------|
| Statistics | <p>Information about frames, bytes, and bits per second received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).</p> <p>The bundle, multilink, and network statistics are reported by the Packet Forwarding Engine (PFE). The Multi Link Detail statistics like fragments, non-fragments and LFI are reported by the PIC.</p> <p>However, the PFE reports an extra overhead of 2 bytes in the output when compared with the Multilink Detail Statistics. This is due to the service-cookie in the PFE which does the link demux for the ML header.</p> <p>The difference in the bytes received and transmitted from Network and Multilink interfaces and Multilink statistics for each member link is divided between the ML and the PPP headers. For example the header counter for a long sequence configuration would be as follows.</p> <ul style="list-style-type: none"> • Input side - Total overhead = 6 bytes. <ul style="list-style-type: none"> • ML: 4 bytes of ML header = 1 byte of Flag + 3 bytes of long sequence number. • PPP: 2 bytes of protocol field. • Output side - Total overhead = 11 bytes. <ul style="list-style-type: none"> • ML: 4 bytes of ML Header = 1 byte of Flag + 3 bytes of Long sequence number. • PPP: 5 bytes = 4 bytes of header + 1 byte of Idle flag. • 2 bytes of Service Cookie. • Bundle—Information for each active bundle link. <ul style="list-style-type: none"> • Multilink: Input and Output—Total number and rate of multilink frames, bytes, and bits per second received and transmitted. It is a module connecting LSQ PIC and its member link. Multilink Input displays L2 fragments received from the member link to the LSQ PIC. Multilink Output displays the L2 fragments transmitted from LSQ PIC to the member links. • Network: Input and Output—Total number of network frames, bytes, and bits per second received and transmitted. It refers to the packets transmitted from an ingress interface to the PFE and then to the LSQ PIC. Network Input displays the L3 packets received from the LSQ PIC to the PFE. Network Output displays the L3 packets transmitted from PFE to LSQ PIC. • Link—Information about links used in the multilink operation. <ul style="list-style-type: none"> • Link name—The interface name of the link services IQ channel and state information (physical link <i>up</i> or <i>down</i>) and up time. • Input and Output—Total number and rate of frames, bytes, and bits per second received and transmitted. | extensive |

Table 119: show interfaces (Link Services IQ) Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------------------|---|------------------------------|
| Multilink detail statistics | <p>Frames, bytes, and bits per second received and sent by the bundle. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).</p> <p>The difference in the bytes received and transmitted from the bundle is divided between the ML and the PPP headers. For example the header counter for a long sequence configuration would be as follows:</p> <ul style="list-style-type: none"> • Input side - Total overhead = 6 bytes. <ul style="list-style-type: none"> • ML: 4 bytes of ML header = 1 byte of Flag + 3 bytes of long sequence number. • PPP: 2 bytes of protocol field. • Output side - Total overhead = 9 bytes. <ul style="list-style-type: none"> • ML: 4 bytes of ML Header = 1 byte of Flag + 3 bytes of Long sequence number. • PPP: 5 bytes = 4 bytes of header + 1 byte of Idle flag. • Bundle—Information for the bundle link. <ul style="list-style-type: none"> • Fragments: Input and Output—Total number and rate of multilink fragments received and transmitted. • Non-fragments: Input and Output—Total number and rate of nonfragmented multilink frames received and transmitted. • LFI: Input and Output—Total number and rate of link fragmented and interleaved frames and bytes. | extensive |
| Protocol | Protocol family configured on the logical interface. | detail extensive none |
| MTU | MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked <i>Adjusted</i> . | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route Table | Routing table in which this address exists. For example, Route table:0 refers to inet.0. | detail extensive |
| Addresses, Flags | Information about the addresses configured on the logical interface. Possible values are described in the "Addresses Flags" section under <i>Common Output Fields Description</i> . | detail extensive none |
| Destination | IP address of the remote side of the connection. | detail extensive none |
| Local | IP address of the logical interface. | detail extensive none |
| Broadcast | Broadcast address on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |

Sample Output

show interfaces extensive (MLPPP on Link Services IQ)

```

user@host> show interfaces lsq-0/2/0 extensive
Physical interface: lsq-0/2/0, Enabled, Physical link is Up
  Interface index: 140, SNMP ifIndex: 25, Generation: 23
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2005-06-02 08:54:36 PDT (00:05:45 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :           8872424           229080 bps
    Output bytes  :           9856960           234448 bps
    Input packets :           38202           117 pps
    Output packets:           39453           117 pps
  Frame exceptions:
    Oversized frames           0
    Errored input frames       0
    Input on disabled link/bundle 0
    Output for disabled link/bundle 0
    Queuing drops              0
  Buffering exceptions:
    Packet data buffer overflow 0
    Fragment data buffer overflow 0
  Assembly exceptions:
    Fragment timeout           0
    Missing sequence number    0
    Out-of-order sequence number 0
    Out-of-range sequence number 0
  Hardware errors (sticky):
    Data memory error          0
    Control memory error       0
  Queue counters:

```

| | Queued packets | Transmitted packets | Dropped packets |
|------|----------------|---------------------|-----------------|
| 0 be | 0 | 0 | 0 |
| 1 ef | 0 | 0 | 0 |
| 2 af | 0 | 0 | 0 |
| 3 nc | 0 | 0 | 0 |

```

  Logical interface lsq-0/2/0.0 (Index 66) (SNMP ifIndex 26) (Generation 5)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
  Bandwidth: 256kbps
  Bundle options:
    MRRU           1504
    Drop timer period 2000
    Sequence number format long (24 bits)
    Fragmentation threshold 0
    Links needed to sustain bundle 1
    Multilink classes 0
    Link layer overhead 4.0 %
  Bundle status:
    Remote MRRU           1500
    Received sequence number 0x0
    Transmit sequence number 0x0
    Packet drops           0 (0 bytes)
    Fragment drops         9 (1401 bytes)

```

```

MRRU exceeded          0
Fragment timeout        0
Missing sequence number 0
Out-of-order sequence number 4
Out-of-range sequence number 0
Packet data buffer overflow 0
Fragment data buffer overflow 0
Statistics              Frames    fps          Bytes        bps
Bundle:
Multilink:
  Input :               79827      239          9593009       232288
  Output:              77533      234          9811743       238056
Network:
  Input :               38202      117          8872424       229080
  Output:              39453      117          9856960       234448
Link:
ds-1/0/2:1:1.0 <-- up
  Input :               1114         87          180183        113608
  Output:              1577        118          199215        119064
ds-1/0/2:1:2.0 <-- down
  Input :               1941        152          187948        118680
  Output:              1574        116          199494        118992
Protocol inet, MTU: 1500 [Adjusted]
Flags: User-MTU, MTU-Protocol-Adjusted
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.74.11/24, Local: 10.74.11.10
Protocol iso, MTU: 1500 [Adjusted]
Flags: User-MTU, MTU-Protocol-Adjusted
Protocol mpls, MTU: 1488 [Adjusted], Maximum labels: 3
Flags: User-MTU, MTU-Protocol-Adjusted

```

show interfaces extensive (Multiclass MLPPP on Link Services IQ)

```

user@host> show interfaces extensive lsq-0/2/0
Physical interface: lsq-0/2/0, Enabled, Physical link is Up
Interface index: 140, SNMP ifIndex: 25, Generation: 23
Link-level type: LinkService, MTU: 1504
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps
Last flapped   : 2005-06-02 08:54:36 PDT (00:02:25 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   :          3474024          223704 bps
Output bytes  :          4193992          233888 bps
Input packets :          15809           116 pps
Output packets:          16788           117 pps
Frame exceptions:
Oversized frames          0
Errored input frames      0
Input on disabled link/bundle 0
Output for disabled link/bundle 0
Queuing drops             0
Buffering exceptions:
Packet data buffer overflow 0
Fragment data buffer overflow 0
Assembly exceptions:
Fragment timeout          0
Missing sequence number   0
Out-of-order sequence number 0
Out-of-range sequence number 0
Hardware errors (sticky):

```

| | | | |
|----------------------|----------------|---------------------|-----------------|
| Data memory error | 0 | | |
| Control memory error | 0 | | |
| Queue counters: | Queued packets | Transmitted packets | Dropped packets |
| 0 be | 0 | 0 | 0 |
| 1 ef | 0 | 0 | 0 |
| 2 af | 0 | 0 | 0 |
| 3 nc | 0 | 0 | 0 |

Logical interface lsq-0/2/0.0 (Index 66) (SNMP ifIndex 26) (Generation 5)

Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP

Bandwidth: 256kbps

Bundle options:

| | |
|--------------------------------|----------------|
| MRRU | 1504 |
| Drop timer period | 2000 |
| Sequence number format | long (24 bits) |
| Fragmentation threshold | 0 |
| Links needed to sustain bundle | 1 |
| Multilink classes | 2 |
| Link layer overhead | 4.0 % |

Multilink class 0 status:

| | |
|-------------------------------|---------------------|
| Received sequence number | 0x4c38 |
| Transmit sequence number | 0x4890 |
| Packet drops | 0 (0 bytes) |
| Fragment drops | 2551 (397084 bytes) |
| MRRU exceeded | 0 |
| Fragment timeout | 52 |
| Missing sequence number | 0 |
| Out-of-order sequence number | 953 |
| Out-of-range sequence number | 0 |
| Packet data buffer overflow | 0 |
| Fragment data buffer overflow | 0 |

Multilink class 1 status:

| | |
|-------------------------------|-------------|
| Received sequence number | 0xffffffff |
| Transmit sequence number | 0x3710 |
| Packet drops | 0 (0 bytes) |
| Fragment drops | 0 (0 bytes) |
| MRRU exceeded | 0 |
| Fragment timeout | 0 |
| Missing sequence number | 0 |
| Out-of-order sequence number | 0 |
| Out-of-range sequence number | 0 |
| Packet data buffer overflow | 0 |
| Fragment data buffer overflow | 0 |

| | | | | |
|------------|--------|-----|-------|-----|
| Statistics | Frames | fps | Bytes | bps |
|------------|--------|-----|-------|-----|

Bundle:

Fragments:

| | | | | |
|---------|-------|-----|---------|--------|
| Input : | 33719 | 239 | 4041763 | 231632 |
| Output: | 32371 | 234 | 4096545 | 237488 |

Packets:

| | | | | |
|---------|-------|-----|---------|--------|
| Input : | 15809 | 116 | 3474024 | 223704 |
| Output: | 16788 | 117 | 4193992 | 233888 |

Multilink class 0:

Fragments:

| | | | | |
|---------|-------|---|---|---|
| Input : | 19331 | 0 | 0 | 0 |
| Output: | 0 | 0 | 0 | 0 |

Packets:

| | | | | |
|---------|------|---|---|---|
| Input : | 2064 | 0 | 0 | 0 |
|---------|------|---|---|---|

```

      Output:          1864          0          0          0
Multilink class 1:
  Fragments:
    Input :           0          0          0          0
    Output:          14096          0          0          0
  Packets:
    Input :          14096          0          0          0
    Output:           0          0          0          0
Link:
  ds-1/0/2:1:1.0, Enabled, Physical link is Up
    Input :           20972          151          2030595          118080
    Output:          16184          116          2048468          118488
  ds-1/0/2:1:2.0, Enabled, Physical link is Up
    Input :           12747           88          2011168          113552
    Output:          16187          118          2048077          119000
Protocol inet, MTU: 1500 [Adjusted], Generation: 14, Route table: 0
Flags: User-MTU, MTU-Protocol-Adjusted
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.0.1.0/30, Local: 10.0.1.2, Broadcast: Unspecified,
  Generation: 18

```

show interfaces extensive (MLPPP on Link Services IQ Bundle)

```

user@host> show interfaces lsq-7/1/0.0 extensive
Logical interface lsq-7/1/0.0 (Index 88) (SNMP ifIndex 114) (Generation 188)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-FR
Last flapped: Never
Bandwidth: 256kbps
Bundle links information:
  Active bundle links      2
  Removed bundle links     0
  Disabled bundle links    0
Bundle options:
  MRRU                      1504
  Drop timer period         1500
  Inner PPP Protocol field compression enabled
  Sequence number format    short (12 bits)
  Fragmentation threshold   0
  Links needed to sustain bundle 1
  Multilink classes         0
  Link layer overhead       4.0 %
Bundle status:
  Received sequence number   0xb74
  Transmit sequence number   0xb74
  Packet drops               0 (0 bytes)
  Fragment drops             0 (0 bytes)
  MRRU exceeded              0
  Fragment timeout           0
  Missing sequence number    0
  Out-of-order sequence number 0
  Out-of-range sequence number 0
  Packet data buffer overflow 0
  Fragment data buffer overflow 0
Statistics      Frames      fps      Bytes      bps
Bundle:
Multilink:
  Input :        315381      0      42757818      0
  Output:        315381      0      43388580      0
Network:
  Input :        315381      0      40952064      0
  Output:        315381      0      40952064      0

```

```

Link:
  ds-6/0/0:1:1.0
    Up time: Up since boot
    Input :      63794      0      25146728      0
    Output:      63778      0      25273164      0
  ds-6/0/0:1:2.0
    Up time: Up since boot
    Input :      251587      0      17611090      0
    Output:      251603      0      18115416      0
Multilink detail statistics:
Bundle:
  Fragments:
    Input :      0      0      0      0
    Output:      0      0      0      0
  Non-fragments:
    Input :      293748      0      19387368      0
    Output:      293748      0      20562360      0
  LFI:
    Input :      21633      0      22152192      0
    Output:      21633      0      22325256      0
Protocol inet, MTU: 1500, Generation: 204, Route table: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.0.1.0/30, Local: 10.0.1.2, Broadcast:
Unspecified, Generation: 214

```

show interfaces extensive (MFR on Link Services IQ Bundle)

```

user@host> show interfaces lsq-1/0/0:0 extensive
Physical interface: lsq-1/0/0:0, Enabled, Physical link is Up
Interface index: 179, SNMP ifIndex: 746, Generation: 182
Link-level type: Multilink-FR-UNI-NNI, MTU: 1508
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
Last flapped   : 2010-11-15 01:11:00 PST (00:31:58 ago)
Statistics last cleared: Never
Hold-times     : Up 0 ms, Down 0 ms
Multilink Frame Relay UNI NNI bundle options:
  Device type      DCE
  MRRU             1508
  Bandwidth        1536kbps
  Fragmentation threshold 0
  Red differential delay limit 120
  Yellow differential delay limit 72
  Red differential delay action Remove link
  Reassembly drop timer 65535
  Links needed to sustain bundle 1
  Link layer overhead 4.0 %
  LIP Hello timer 10
    Acknowledgement timer 4
    Acknowledgement retries 2
  Bundle class     A
  LMI type         Consortium
    T391 LIV polling timer 10
    T392 polling verification timer 15
    N391 full status polling count 6
    N392 error threshold 3
    N393 monitored event count 4
  Consortium LMI settings: n392dce 3, n393dce 4, t392dce 15 seconds
LMI statistics:
  Input : 188 (last seen 00:00:01 ago)

```



```

Output: 189 (last sent 00:00:01 ago)
DTE statistics:
  Enquiries sent : 0
  Full enquiries sent : 0
  Enquiry responses received : 0
  Full enquiry responses received : 0
DCE statistics:
  Enquiries received : 157
  Full enquiries received : 31
  Enquiry responses sent : 158
  Full enquiry responses sent : 31
Common statistics:
  Unknown messages received : 0
  Asynchronous updates received : 0
  Out-of-sequence packets received : 0
  Keepalive responses timedout : 0
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Multilink Frame Relay UNI NNI bundle errors:
  Packet drops 0 (0 bytes)
  Fragment drops 0 (0 bytes)
  MRRU exceeded 0
  Exception events 0
Multilink Frame Relay UNI NNI bundle statistics:
      Frames      fps      Bytes      bps

Multilink:
  Input : 0 0 0 0
  Output: 0 0 0 0
Network:
  Input : 0 0 0 0
  Output: 0 0 0 0
Multilink Frame Relay UNI NNI bundle links information:
  Active bundle links 1
  Removed bundle links 0
  Disabled bundle links 0
Multilink Frame Relay UNI NNI active bundle links statistics:
      Frames      fps      Bytes      bps

t1-7/0/0:1:3.0
Up time: 00:31:24
  Input : 0 0 0 0
  Output: 0 0 0 0
  Current differential delay 0.0 ms
  Recent high differential delay 0.0 ms
  Times over red diff delay 0
  Times over yellow diff delay 0
LIP:add_lnk lnk_ack lnk_rej hello hel_ack lnk_rem rem_ack
Rcv: 2 2 0 0 189 0 0
Xmt: 2 1 0 189 0 0 0

```

Logical interface lsq-1/0/0:2.0 (Index 77) (SNMP ifIndex 751) (Generation 142)

Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI

```

Last flapped: 2010-11-15 01:11:40 PST (00:31:18 ago)
Bundle status:
  Received sequence number      0xfff
  Transmit sequence number      0x0
  Packet drops                  0 (0 bytes)
  Fragment drops                0 (0 bytes)
  MRRU exceeded                 0
  Fragment timeout              0
  Missing sequence number       0
  Out-of-order sequence number  0
  Out-of-range sequence number  0
  Packet data buffer overflow   0
  Fragment data buffer overflow 0
Statistics      Frames      fps      Bytes      bps
Bundle:
Multilink:
  Input :         0         0         0         0
  Output:         0         0         0         0
Network:
  Input :         0         0         0         0
  Output:         0         0         0         0
Link:
  t1-7/0/0:1:3.0
  Up time: 00:31:24
  Input :         0         0         0         0
  Output:         0         0         0         0
Multilink detail statistics:
Bundle:
Fragments:
  Input :         0         0         0         0
  Output:         0         0         0         0
Non-fragments:
  Input :         0         0         0         0
  Output:         0         0         0         0
Protocol inet, MTU: 1500, Generation: 153, Route table: 0
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.0.1.8/30, Local: 10.0.1.9, Broadcast: Unspecified,
Generation: 154
DLCI 12
Flags: Active
Total down time: 00:00:32 sec, Last down: 00:31:50 ago
Traffic statistics:
  Input bytes :           0
  Output bytes :          0
  Input packets:          0
  Output packets:         0
DLCI statistics:
  Active DLCI :1 Inactive DLCI :0

```

show interfaces (Multiclass MLPPP on Link Services IQ)

```

user@host> show interfaces extensive lsq-0/2/0
Physical interface: lsq-0/2/0, Enabled, Physical link is Up
Interface index: 140, SNMP ifIndex: 25, Generation: 23
Link-level type: LinkService, MTU: 1504
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps
Last flapped : 2005-06-02 08:54:36 PDT (00:02:25 ago)
Statistics last cleared: Never
Traffic statistics:

```

```

Input bytes :          3474024          223704 bps
Output bytes :         4193992          233888 bps
Input packets:         15809           116 pps
Output packets:        16788           117 pps
Frame exceptions:
  Oversized frames      0
  Errored input frames  0
  Input on disabled link/bundle 0
  Output for disabled link/bundle 0
  Queuing drops        0
Buffering exceptions:
  Packet data buffer overflow 0
  Fragment data buffer overflow 0
Assembly exceptions:
  Fragment timeout      0
  Missing sequence number 0
  Out-of-order sequence number 0
  Out-of-range sequence number 0
Hardware errors (sticky):
  Data memory error     0
  Control memory error  0
Queue counters:         Queued packets  Transmitted packets  Dropped packets

0 be                    0                0                0
1 ef                    0                0                0
2 af                    0                0                0
3 nc                    0                0                0

```

Logical interface lsq-0/2/0.0 (Index 66) (SNMP ifIndex 26) (Generation 5)

Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP

Bandwidth: 256kbps

Bundle options:

```

MRRU                    1504
Drop timer period      2000
Sequence number format long (24 bits)
Fragmentation threshold 0
Links needed to sustain bundle 1
Multilink classes      2
Link layer overhead    4.0 %

```

Multilink class 0 status:

```

Received sequence number 0x4c38
Transmit sequence number 0x4890
Packet drops             0 (0 bytes)
Fragment drops           2551 (397084 bytes)
MRRU exceeded            0
Fragment timeout         52
Missing sequence number  0
Out-of-order sequence number 953
Out-of-range sequence number 0
Packet data buffer overflow 0
Fragment data buffer overflow 0

```

Multilink class 1 status:

```

Received sequence number 0xffffffff
Transmit sequence number 0x3710
Packet drops             0 (0 bytes)
Fragment drops           0 (0 bytes)
MRRU exceeded            0
Fragment timeout         0

```

```

Missing sequence number      0
Out-of-order sequence number 0
Out-of-range sequence number 0
Packet data buffer overflow  0
Fragment data buffer overflow 0
Statistics      Frames      fps      Bytes      bps
Bundle:
Fragments:
  Input :      33719      239      4041763      231632
  Output:      32371      234      4096545      237488
Packets:
  Input :      15809      116      3474024      223704
  Output:      16788      117      4193992      233888
Multilink class 0:
Fragments:
  Input :      19331      0      0      0
  Output:      0      0      0      0
Packets:
  Input :      2064      0      0      0
  Output:      1864      0      0      0
Multilink class 1:
Fragments:
  Input :      0      0      0      0
  Output:      14096      0      0      0
Packets:
  Input :      14096      0      0      0
  Output:      0      0      0      0
Link:
ds-1/0/2:1:1.0, Enabled, Physical link is Up
  Input :      20972      151      2030595      118080
  Output:      16184      116      2048468      118488
ds-1/0/2:1:2.0, Enabled, Physical link is Up
  Input :      12747      88      2011168      113552
  Output:      16187      118      2048077      119000
Protocol inet, MTU: 1500 [Adjusted], Generation: 14, Route table: 0
Flags: User-MTU, MTU-Protocol-Adjusted
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.0.1.0/30, Local: 10.0.1.2, Broadcast: Unspecified,
  Generation: 18

```

show interfaces (Multilink Services)

| | |
|---------------------------------|---|
| Syntax | <pre>show interfaces ml-fpc/pic/port <brief detail extensive terse> <descriptions> <media> <snmp-index> <statistics></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M Series and T Series routers only) Display status information about the specified multilink services interface. |
| Options | <p>ml-fpc/pic/port—Display standard status information about the specified multilink services interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index—(Optional) Display the SNMP index of interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| List of Sample Output | show interfaces extensive (Multilink Services) on page 2161 |
| Output Fields | Table 120 on page 2155 lists the output fields for the show interfaces (Multilink Services) command. Output fields are listed in the approximate order in which they appear. |

Table 120: Multilink Services show interfaces Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|--|---------------------------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface's index number, which reflects its initialization sequence. | detail extensive
none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive
none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |

Table 120: Multilink Services show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|---------------------------------|
| Link-level type | Encapsulation being used on the physical interface: Multilink . | All levels |
| MTU | MTU size on the physical interface. | All levels |
| Device flags | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) . | detail extensive
none |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Traffic statistics | Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. | detail extensive |
| Frame exceptions | Information about framing exceptions. Includes events recorded under Exception Events for each logical interface: <ul style="list-style-type: none"> • Oversized frames—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb (kilobits). • Errored input frames—Number of input frame errors. • Input on disabled link/bundle—Number of frames received on disabled links. These can result either from an inconsistent configuration, or from a bundle or link being brought up or down with traffic actively flowing through it. • Output for disabled link/bundle—Number of frames sent for a disabled or unavailable link. These can result either from an inconsistent configuration, or from a bundle being brought up or down with traffic actively flowing through it. • Queuing drops—Total number of packets dropped before traffic enters the link services IQ interface. Indicates that the interface is becoming oversubscribed. | extensive |
| Buffering exceptions | Information about buffering exceptions. Includes events recorded under Exception Events for each logical interface: <ul style="list-style-type: none"> • Packet data buffer overflow—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical multilink services interface capacity. • Fragment data buffer overflow—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical multilink services interface capacity. Check the logical interface exception event counters to determine which bundle is responsible. | extensive |

Table 120: Multilink Services show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------------|---|------------------|
| Assembly exceptions | <p>Information about assembly exceptions. Includes events recorded under Exception Events for each logical interface.</p> <p>An assembly exception does not necessarily indicate an operational problem with the Multilink PIC itself. If multilink-encapsulated traffic is dropped or reordered after a sequence number has been assigned, the assembling multilink interface records one or more exception events. The multilink interface can drop multilink-encapsulated fragments itself as a result. Any multilink packets or fragments dropped by the Multilink PIC itself result in packet or fragment drop counts on individual logical interfaces. If the logical interface drop counts are zero, but exception events are seen, the most likely cause is a problem with the individual link interfaces. Even if the logical interface fragment drop counts are nonzero, excess differential delay or traffic losses on individual interfaces can be the root cause.</p> <ul style="list-style-type: none"> • Fragment timeout—Drop-timer expired while a fragment sequence number was outstanding. Occurs only if drop-timer is enabled. This can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. These events do not necessarily indicate any problem with the operation of the Multilink PIC itself. If one or more individual links drop traffic, these events can occur. Check the logical interface exception event counters to determine which bundle is responsible. • Missing sequence number—A gap was detected in the sequence numbers of fragments on a bundle. These events do not necessarily indicate any problem with the operation of the Multilink PIC itself. If one or more individual links drop traffic, these events can occur. Check the logical interface exception event counters to determine which bundle is responsible. • Out-of-order sequence number—Two frames with out-of-order sequence numbers occurred within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. Check the logical interface exception event counters to determine which bundle is responsible. • Out-of-range sequence number—Frame was received with out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost, or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. Check the logical interface exception event counters to determine which bundle is responsible. | extensive |
| Hardware errors | <p>Information about hardware errors:</p> <ul style="list-style-type: none"> • Data memory error—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. • Control memory error—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. | extensive |
| Logical Interface | | |
| Logical interface | Logical interface name. | All levels |

Table 120: Multilink Services show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------|---|---------------------------------|
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive
none |
| SNMP ifIndex | SNMP interface index number. | detail extensive
none |
| Encapsulation | Encapsulation being used: PPP or Multilink PPP. | All levels |
| Bandwidth | Speed at which the interface is running. | All levels |
| Flags | Logical interface flags. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> . | detail extensive
none |
| Bundle options | Information about configured bundle options: <ul style="list-style-type: none"> • MRRU—Configured size of the MRRU (maximum received reconstructed unit). It can be 1500 to 4500 bytes. • Drop timer period—Configured drop timeout period. It can be 0 through 127 ms. A value of 0 disables the timer. The default setting is 0. • Sequence number format—Configured size of the sequence header: 12 or 24 bits. The default is 24 bits. • Fragmentation threshold—Configured fragmentation threshold. A value of 0 results in no fragmentation. Nonzero values can be 128 through 16,320 bytes, in integer multiples of 64 bytes. The default is 0. • Links needed to sustain bundle—Minimum number of links to sustain the bundle: 1 through 8. | detail extensive
none |

Table 120: Multilink Services show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|--|---------------------------------|
| Bundle status (MLPPP) or Multilink class status (MC-MLPPP) | <p>Information about bundle status:</p> <ul style="list-style-type: none"> • Remote MRRU—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed. • Received sequence number—Sequence number for received packets. • Transmitted sequence number—Sequence number for transmitted packets. • Packet drops—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • Fragment drops—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • MRRU exceeded—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • Fragment timeout—Drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • Missing sequence number—Gap detected in the sequence numbers of fragments on a bundle. • Out-of-order sequence number—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • Out-of-range sequence number—Frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • Packet data buffer overflow—Packet buffer memory full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • Fragment data buffer overflow—Fragment buffer memory full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. | detail extensive |
| Remote MRRU | MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed. | detail extensive
none |

Table 120: Multilink Services show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|--|---------------------------------|
| Bundle errors | Information about bundle errors: <ul style="list-style-type: none"> • Packet drops—Number and byte count of output packets dropped, rather than being encapsulated and sent out of the router as fragments. • Fragment drops—Number and byte count of input fragments dropped, rather than being reassembled and handled by the router as packets. • MRRU exceeded—Number of reassembled packets exceeding the MRRU. • Exception events—Number of exceptional events encountered while handling traffic on the bundle, other than MRRU exceeded errors. These events are categorized under the physical interface: Frame exceptions, Buffering exceptions, and Fragment exceptions. Exception events do not necessarily indicate that the multilink interface is not operating properly. Individual link failures can produce exceptional events. | detail extensive |
| Statistics | Information about fragments and packets received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. <ul style="list-style-type: none"> • Bundle—Information about bundles. • Link—Information about links used in the multilink operation. | detail extensive |
| Protocol | Protocol family configured on the logical interface. | detail extensive
none |
| MTU | MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked Adjusted . | detail extensive
none |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive
none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route table | Route table in which this address exists. For example, Route table:0 refers to inet.0. | detail extensive |
| Addresses, Flags | Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . | detail extensive
none |
| Destination | IP address of the remote side of the connection. | detail extensive
none |
| Local | IP address of the logical interface. | detail extensive
none |
| Broadcast | Broadcast address on the logical interface. | detail extensive
none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |

Sample Output

show interfaces extensive (Multilink Services)

```

user@host> show interfaces ml-0/3/0 extensive
Physical interface: ml-0/3/0, Enabled, Physical link is Up
  Interface index: 273, SNMP ifIndex: 196, Generation: 535
  Link-level type: Multilink, MTU: 4474
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2002-04-25 14:21:34 PDT (21:06:59 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :                3535                0 bps
    Output bytes  :                4135                0 bps
    Input packets :                 87                0 pps
    Output packets:                103                0 pps
  Frame exceptions:
    Oversized frames          0
    Errored input frames      0
    Input on disabled link/bundle 0
    Output for disabled link/bundle 0
    Queuing drops             0
  Buffering exceptions:
    Packet data buffer overflow 0
    Fragment data buffer overflow 0
  Assembly exceptions:
    Fragment timeout           0
    Missing sequence number    0
    Out-of-order sequence number 0
    Out-of-range sequence number 0
  Hardware errors (sticky):
    Data memory error          0
    Control memory error       0

  Logical interface ml-0/3/0.1 (Index 110) (SNMP ifIndex 674)
  (Generation 402)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
    Bandwidth: 12288kbps
    Bundle options:
      MRRU                1524
      Drop timer period    0
      Sequence number format long (24 bits)
      Fragmentation threshold 0
      Links needed to sustain bundle 1
    Bundle status:
      Remote MRRU          1500
      Received sequence number 0x19ec14
      Transmit sequence number 0x38cfa8
      Packet drops          0 (0 bytes)
      Fragment drops        0 (0 bytes)
      MRRU exceeded         0
      Fragment timeout       0
      Missing sequence number 0
      Out-of-order sequence number 0
      Out-of-range sequence number 0
      Packet data buffer overflow 0
      Fragment data buffer overflow 0
    Bundle errors:
      Packet drops          0 (0 bytes)
      Fragment drops        0 (0 bytes)

```

```

MRRU exceeded          0
Exception events        0
Statistics              Frames      fps      Bytes      bps
Bundle:
Fragments:
  Input :               5          0        450        0
  Output:              6          0        499        0
Packets:
  Input :               5          0        450        0
  Output:             12          0       1202        0
Link:
t1-0/1/0:11.0
  Input :               1          0         90        0
  Output:              1          0         92        0
t1-0/1/0:12.0
  Input :               1          0         90        0
  Output:              1          0         92        0
t1-0/1/0:10.0
  Input :               1          0         90        0
  Output:              1          0         92        0
t1-0/1/0:14.0
  Input :               1          0         90        0
  Output:              1          0         92        0
t1-0/1/0:13.0
  Input :               1          0         90        0
  Output:              1          0         92        0
t1-0/1/0:8.0
  Input :               0          0          0        0
  Output:              0          0          0        0
t1-0/1/0:9.0
  Input :               0          0          0        0
  Output:              0          0          0        0
Protocol inet, MTU: 1500 [Adjusted], Flags: Generation: 752 Route table: 0
Addresses, Flags: Is-Preferred Is-Primary, MTU-Protocol-Adjusted
Destination: 1.1.2.2, Local: 1.1.2.1, Broadcast: Unspecified,
Generation: 1090
Protocol iso, MTU: 1500 [Adjusted], Flags: Is-Primary,
Generation: 753 Route table: 0

```

Monitoring, Sampling, and Collection Services Operational Commands

- `clear passive-monitoring statistics`
- `clear services accounting statistics inline-jflow`
- `clear services dynamic-flow-capture`
- `clear services flow-collector statistics`
- `clear services rpm twamp server connection`
- `clear services video-monitoring mdi errors fpc-slot`
- `clear services video-monitoring mdi statistics fpc-slot`
- `request services flow-collector change-destination primary interface`
- `request services flow-collector change-destination secondary interface`
- `request services flow-collector test-file-transfer`
- `show forwarding-options next-hop-group`
- `show forwarding-options port-mirroring`

- `show interfaces (Dynamic Flow Capture)`
- `show interfaces (Flow Collector)`
- `show interfaces (Flow Monitoring)`
- `show passive-monitoring error`
- `show passive-monitoring flow`
- `show passive-monitoring memory`
- `show passive-monitoring status`
- `show passive-monitoring usage`
- `show services accounting aggregation`
- `show services accounting aggregation template`
- `show services accounting errors`
- `show services accounting flow`
- `show services accounting flow-detail`
- `show services accounting memory`
- `show services accounting packet-size-distribution`
- `show services accounting status`
- `show services accounting usage`
- `show services dynamic-flow-capture content-destination`
- `show services dynamic-flow-capture control-source`
- `show services dynamic-flow-capture statistics`
- `show services flow-collector file interface`
- `show services flow-collector input interface`
- `show services flow-collector interface`
- `show services rpm active-servers`
- `show services rpm history-results`
- `show services rpm probe-results`
- `show services rpm rfc2544-benchmarking`
- `show services rpm rfc2544-benchmarking test-id`
- `show services rpm twamp server connection`
- `show services rpm twamp server session`
- `show services video-monitoring mdi errors fpc-slot`
- `show services video-monitoring mdi flows fpc-slot`
- `show services video-monitoring mdi stats fpc-slot`
- `test services rpm rfc2544-benchmarking test`

clear passive-monitoring statistics

| | |
|---------------------------------|---|
| Syntax | clear passive-monitoring statistics (all interface <i>interface-name</i>) |
| Release Information | Command introduced in Junos OS Release 7.6. |
| Description | (M40e, M160, and M320 routers and T Series routers only) Clear statistics for one passive monitoring interface or for all passive monitoring interfaces. |
| Options | all —Clear statistics for all configured passive monitoring interfaces.

interface <i>interface-name</i> —Clear statistics for the specified passive monitoring interface (<i>mo-fpc/pic/port</i>). |
| Required Privilege Level | network |
| List of Sample Output | clear passive-monitoring statistics on page 2164 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear passive-monitoring statistics

```
user@host> clear passive-monitoring statistics interface mo-5/0/0
```

clear services accounting statistics inline-jflow

| | |
|---------------------------------|---|
| Syntax | clear services accounting statistics inline-jflow
<inline-jflow (fpc-slot <i>slot-number</i>)> |
| Release Information | Command introduced in Junos OS Release 14.2 for MX Series routers. |
| Description | Clear inline flow statistics for a specified FPC. |
| Options | <p>fpc-slot <i>slot-number</i>—Clear inline flow statistics for the specified FPC.</p> <ul style="list-style-type: none"> • MX80 routers only—Replace <i>slot-number</i> with a value from 0 through 1. • MX104 routers only—Replace <i>slot-number</i> with a value from 0 through 2. • MX240 routers only—Replace <i>slot-number</i> with a value from 0 through 2. • MX480 routers only—Replace <i>slot-number</i> with a value from 0 through 5. • MX960 routers only—Replace <i>slot-number</i> with a value from 0 through 11. • MX2010 routers only—Replace <i>slot-number</i> with a value from 0 through 9. • MX2020 routers only—Replace <i>slot-number</i> with a value from 0 through 19. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show services accounting flow on page 2213 |

Sample Output

clear services accounting statistics inline-jflow

```
user@host> regress@mobsoln480b# run clear services accounting statistics inline-jflow fpc-slot
5
Statistics Cleared
```

clear services dynamic-flow-capture

| | |
|---------------------------------|---|
| Syntax | <code>clear services dynamic-flow-capture capture-group <i>group-name</i></code>
<code><criteria-identifier <i>identifier</i>></code>
<code><destination-identifier <i>identifier</i>></code>
<code><force></code>
<code><static></code> |
| Release Information | Command introduced in Junos OS Release 7.4. |
| Description | (M320 routers and T Series routers only) Clear dynamic flow capture information for specified capture group. |
| Options | <code>capture-group <i>group-name</i></code> —Capture-group identifier.

<code>criteria-identifier <i>identifier</i></code> —(Optional) Criteria identifier.

<code>destination-identifier <i>identifier</i></code> —(Optional) Content destination identifier.

<code>force</code> —(Optional) Force clearing of criteria.

<code>static</code> —(Optional) Clear static criteria. |
| Required Privilege Level | network |
| List of Sample Output | clear services dynamic-flow-capture on page 2166 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services dynamic-flow-capture

```
user@host> clear services dynamic-flow-capture capture-group flow-a
```


clear services flow-collector statistics

| | |
|---------------------------------|--|
| Syntax | clear services flow-collector statistics (all interface <i>interface-name</i>) |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M40e, M160, and M320 routers and T Series routers only) Clear statistics for one flow collector interface or for all flow collector interfaces. |
| Options | <p>all—Clear statistics for all configured flow collector interfaces.</p> <p>interface <i>interface-name</i>—Clear statistics for the specified flow collector interface (<i>cp-fpc/pic/port</i>).</p> |
| Required Privilege Level | network |
| List of Sample Output | clear services flow-collector statistics on page 2167 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear services flow-collector statistics

```

user@host> clear services flow-collector statistics interface cp-5/0/0
Flow collector interface: cp-5/0/0
Interface state: Collecting flows
Statistics cleared successfully

```

clear services rpm twamp server connection

| | |
|---------------------------------|--|
| Syntax | clear services rpm twamp server connection
<i><connection-id></i> |
| Release Information | Command introduced in Junos OS Release 9.3. |
| Description | Clear connections established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default all established connections are cleared (along with the sessions on those connections). To clear only a specific connection, specify the connection ID when you issue the command. |
| Options | <i>connection-id</i> —(Optional) Clear only the specified connection. |
| Required Privilege Level | clear |

clear services video-monitoring mdi errors fpc-slot

| | |
|---------------------------------|--|
| Syntax | clear services video-monitoring mdi errors fpc-slot <i>fpc-slot</i> |
| Release Information | Command introduced in Junos OS Release 14.1. |
| Description | Clear all media delivery index error counters. |
| Options | fpc-slot —Number of the fpc slot. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show services video-monitoring mdi stats fpc-slot on page 2291 |

clear services video-monitoring mdi statistics fpc-slot

| | |
|---------------------------------|--|
| Syntax | clear services video-monitoring mdi statistics fpc-slot <i>fpc-slot</i> |
| Release Information | Command introduced in Junos OS Release 14.1. |
| Description | Clear all media delivery index statistics counters except for active flows. |
| Options | fpc-slot —Number of the fpc slot. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show services video-monitoring mdi stats fpc-slot on page 2291 |

request services flow-collector change-destination primary interface

| | |
|---------------------------------|--|
| Syntax | request services flow-collector change-destination primary interface <i>cp-fpc/pic/port</i>
<clear-files>
<clear-logs>
<immediately gracefully> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M40e, M160, and M320 routers and T Series routers only) Switch to the primary File Transfer Protocol (FTP) server that is configured as a flow collector. |
| Options | <p>none—Switch to the primary FTP server.</p> <p>cp-fpc/pic/port—Specify the flow collector interface name for the primary destination.</p> <p>clear-files—(Optional) Request clearing of existing data files in the FTP wait queue when the switch takes place.</p> <p>clear-logs—(Optional) Request clearing of existing logs when the switch takes place.</p> <p>immediately gracefully—(Optional) Specify whether you want the switch to take place immediately, or to affect only newly created files.</p> |
| Required Privilege Level | maintenance |
| List of Sample Output | request services flow-collector change-destination primary interface on page 2171 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request services flow-collector change-destination primary interface

```

user@host> request services flow-collector change-destination primary interface cp-6/0/0
Flow collector interface: cp-6/0/0
Interface state: Collecting flows
Destination change successful

```

request services flow-collector change-destination secondary interface

| | |
|---------------------------------|---|
| Syntax | <code>request services flow-collector change-destination secondary interface <i>cp-fpc/pic/port</i></code>
<code><clear-files></code>
<code><clear-logs></code>
<code><immediately gracefully></code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M40e, M160, and M320 routers and T Series routers only) Switch to the secondary File Transfer Protocol (FTP) server that is configured as a flow collector. |
| Options | <p>none—Switch to the secondary FTP server.</p> <p><i>cp-fpc/pic/port</i>—Specify the flow collector interface name (<i>cp-fpc/pic/port</i>) for the secondary destination.</p> <p>clear-files—(Optional) Request clearing of existing data files in the FTP wait queue when the switch takes place.</p> <p>clear-logs—(Optional) Request clearing of existing logs when the switch takes place.</p> <p>immediately gracefully—(Optional) Specify whether you want the switch to take place immediately, or to affect only newly created files.</p> |
| Required Privilege Level | maintenance |
| List of Sample Output | request services flow-collector change-destination secondary interface on page 2172 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request services flow-collector change-destination secondary interface

```
user@host> request services flow-collector change-destination secondary interface cp-6/0/0
Flow collector interface: cp-6/0/0
Interface state: Collecting flows
Destination change successful
```

request services flow-collector test-file-transfer

| | |
|---------------------------------|--|
| Syntax | <code>request services flow-collector test-file-transfer <i>filename</i> interface (all <i>cp-fpc/pic/port</i>) (channel-zero channel-one) (primary secondary)</code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M40e, M160, and M320 routers, PTX Series, and T Series routers only) Transfer a test file to the primary or secondary File Transfer Protocol (FTP) server that is configured as a flow collector. This command verifies that the output side of the flow collector interface is operating properly. |
| Options | <p><i>filename</i>—Name of the test file to transfer.</p> <p>interface all <i>cp-fpc/pic/port</i>—Transfer a test file of flows from all configured flow collector interfaces or from only the specified interface.</p> <p>channel-zero channel-one—Transfer a file from export channel 0 (unit 0) or channel 1 (unit 1) of the PIC.</p> <p>primary secondary—Transfer a file to the primary or secondary server configured as a flow collector.</p> |
| Required Privilege Level | network |
| List of Sample Output | request services flow-collector test-file-transfer on page 2173 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request services flow-collector test-file-transfer

```
user@router> request services flow-collector test-file-transfer test_file interface cp-7/1/0
channel-one primary
```

```
Flow collector interface: cp-7/1/0
Interface state: Collecting flows
Response: Test file transfer successfully scheduled
```

show forwarding-options next-hop-group

| | |
|---------------------------------|--|
| Syntax | show forwarding-options next-hop-group
<terse brief detail>
<group-name> |
| Release Information | Command introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Support for IPv6 introduced in Junos OS Release 14.2 for the MX Series routers. |
| Description | Display current state of next-hop groups. |
| Options | terse brief detail —(Optional) Display the specified level of output.

group-name —(Optional) Display a single next-hop group. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show forwarding-options port-mirroring on page 2177 |
| List of Sample Output | show forwarding-options next-hop-group terse on page 2175
show forwarding-options next-hop-group brief on page 2175
show forwarding-options next-hop-group detail on page 2175 |
| Output Fields | Table 121 on page 2174 lists the output fields for the show forwarding-options next-hop-group command. Output fields are listed in the approximate order in which they appear. |

Table 121: show forwarding-options next-hop-group Output Fields

| Field Name | Field Description | Level of Output |
|--------------------------------------|---|---------------------|
| Next-hop-group | Name of next-hop group. | All levels |
| Type | Next-hop group type, such as inet , inet6 or layer-2 . | All levels |
| State | Next-hop group state, either up or down . | All levels |
| Members Interfaces | Names of interfaces to which next-hop group members belong. | brief detail |
| Member Subgroup | Names of subgroups to which next-hop group members belong. | brief detail |
| Number of members configured | Number of next-hop group members configured. | detail |
| Number of members that are up | Number of next-hop group members that are up. | detail |

Table 121: show forwarding-options next-hop-group Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|----------------------------------|-----------------|
| Number of subgroups configured | Number of subgroups configured. | detail |
| Number of subgroups that are up | Number of subgroups that are up. | detail |

Sample Output

show forwarding-options next-hop-group terse

```

user@host> show forwarding-options next-hop-group terse
Next-hop-group      Type      State
nhg                  inet      up
nhg6                 inet6     up
vpls_nhg_2          layer-2   down

```

show forwarding-options next-hop-group brief

```

user@host> show forwarding-options next-hop-group brief

Next-hop-group: nhg
Type: inet
State: up
Members Interfaces:
  ge-0/2/8.0      next-hop 30.1.1.10
  ge-5/1/8.0      next-hop 10.1.1.10
  ge-5/1/9.0      next-hop 20.1.1.10

Next-hop-group: nhg6
Type: inet6
State: up
Members Interfaces:
  ge-5/1/5.0      next-hop 10::1:1:10
  ge-5/1/6.0      next-hop 20::1:1:10
Member Subgroup: nhsg6
Members Interfaces:
  ge-5/0/4.0      next-hop 3::1:1:1
  ge-5/1/4.0      next-hop 4::1:1:1

Next-hop-group: vpls_nhg_2
Type: layer-2      State: down

```

show forwarding-options next-hop-group detail

```

user@host> show forwarding-options next-hop-group detail

Next-hop-group: nhg
Type: inet
State: up
Number of members configured      : 3
Number of members that are up    : 3
Number of subgroups configured    : 0
Number of subgroups that are up  : 0

```

| Members Interfaces: | | State |
|---------------------|--------------------|-------|
| ge-0/2/8.0 | next-hop 30.1.1.10 | up |
| ge-5/1/8.0 | next-hop 10.1.1.10 | up |
| ge-5/1/9.0 | next-hop 20.1.1.10 | up |

Next-hop-group: nhg6

Type: inet6

State: up

Number of members configured : 2

Number of members that are up : 2

Number of subgroups configured : 1

Number of subgroups that are up : 1

| Members Interfaces: | | State |
|---------------------|---------------------|-------|
| ge-5/1/5.0 | next-hop 10::1:1:10 | up |
| ge-5/1/6.0 | next-hop 20::1:1:10 | up |

Member Subgroup: nhsg6

Number of members configured : 2

Number of members that are up : 2

| Members Interfaces: | | State |
|---------------------|-------------------|-------|
| ge-5/0/4.0 | next-hop 3::1:1:1 | up |
| ge-5/1/4.0 | next-hop 4::1:1:1 | up |

Next-hop-group: vpls_nhg_2

Number of members configured : 2

Number of members that are up : 0

Number of subgroups configured : 0

Number of subgroups that are up : 0

Type: layer-2 State: down

Members Interfaces: State

ge-2/2/1.100 down

ge-2/3/9.0 down

show forwarding-options port-mirroring

| | |
|---------------------------------|---|
| Syntax | show forwarding-options port-mirroring
<terse detail>
<instance-name> |
| Release Information | Command introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches. |
| Description | Display current state of port-mirroring instances. |
| Options | terse detail —(Optional) Display the specified level of output.

instance-name —(Optional) Display a single port-mirroring instance. |
| Required Privilege Level | view |
| Related Documentation | |
| List of Sample Output | show forwarding-options port-mirroring terse on page 2178
show forwarding-options port-mirroring detail on page 2178 |
| Output Fields | Table 122 on page 2177 lists the output fields for the show forwarding-options port-mirroring command. Output fields are listed in the approximate order in which they appear. |

Table 122: show forwarding-options port-mirroring Output Fields

| Field Name | Field Description | Level of Output |
|--------------------------|---|-----------------|
| Instance Name | Name of port-mirroring instance. | All levels |
| Instance Id | Instance identification number. | All levels |
| State | Instance state, either up or down . | All levels |
| Input parameters | | |
| Rate | Rate (ratio of packets sampled). | detail |
| Run-length | Run length (number of consecutive packets sampled). | detail |
| Maximum-packet-length | Maximum packet length. | detail |
| Output parameters | | |
| Family | Protocol family. | detail |
| State | Instance state, either up or down . | detail |
| Destination | Destination (next-hop group name). | detail |

Sample Output

show forwarding-options port-mirroring terse

```
user@host> show forwarding-options port-mirroring terse
Instance Name      Instance Id  State
&global_instance   1           up
inst1              2           up
```

show forwarding-options port-mirroring detail

```
user@host> show forwarding-options port-mirroring detail
Instance Name: &global_instance
Instance Id: 1      State: up
  Input parameters:
    Rate:          10
    Run-length:     4
    Maximum-packet-length: 0
  Output parameters:
    Family: inet    State: up Destination: inet_nhg
    Family: vpls/eth-switch State: up Destination: vpls_nhg

Instance Name: inst1
Instance Id: 2      State: up
  Input parameters:
    Rate:          1
    Run-length:     0
    Maximum-packet-length: 200
  Output parameters:
    Family: inet    State: up Destination: inet_nhg
    Family: vpls/eth-switch State: down Destination: vpls_nhg_2
```

show interfaces (Dynamic Flow Capture)

| | |
|---------------------------------|---|
| Syntax | <pre>show interfaces dfc-<i>fpc/pic/port:channel</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | Command introduced in Junos OS Release 7.4. |
| Description | (M320 and M120 routers and T Series routers only) Display status information about the specified dynamic flow capture interface. |
| Options | <p>dfc-<i>fpc/pic/port:channel</i>—Display standard status information about the specified dynamic flow capture interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| List of Sample Output | show interfaces (Dynamic Flow Capture) on page 2182 |
| Output Fields | Table 123 on page 2179 lists the output fields for the show interfaces (Dynamic Flow Capture) command. Output fields are listed in the approximate order in which they appear. |

Table 123: Dynamic Flow Capture show interfaces Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|--|------------------------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |
| Type | Type of interface. | All levels |

Table 123: Dynamic Flow Capture show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|--|-------------------------|
| Link-level type | Encapsulation type used on the physical interface. | All levels |
| MTU | Maximum Transmit Unit (MTU). Size of the largest packet to be transmitted. | All levels |
| Speed | Network speed on the interface. | All levels |
| Device flags | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Link type | Data transmission type. | All levels |
| Link flags | Information about the link. Possible values are described in the “Link Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour:minute:second timezone (hour:minute:second ago)</i> . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) . | detail extensive |
| Input Rate | Input rate in bits per second (bps) and packets per second (pps). | None specified |
| Output Rate | Output rate in bps and pps. | None specified |
| Traffic statistics | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> Input rate, Output rate—Number of bits per second (packets per second) received and transmitted on the interface. Input packets, Output packets—Number of packets received and transmitted on the interface. | detail extensive |
| Input errors | <ul style="list-style-type: none"> Errors—Input errors on the interface. Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. Framing errors—Number of packets received with an invalid frame checksum (FCS). Runts—Frames received smaller than the runt threshold. Giants—Frames received larger than the giant threshold. Policed Discards—Frames that the incoming packet match code discarded because the frames did not recognize them or were not of interest. Usually, this field reports protocols that the Junos OS does not support. Resource errors—Sum of transmit drops. | extensive |

Table 123: Dynamic Flow Capture show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|---|------------------------------|
| Output errors | <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • Errors—Sum of outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet dropped by the ASIC RED mechanism. • Resource errors—Sum of transmit drops. | extensive |
| Logical Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Flags | Information about the logical interface; values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Encapsulation | Encapsulation on the logical interface. | All levels |
| Input packets | Number of packets received on the logical interface. | None specified |
| Output packets | Number of packets transmitted on the logical interface. | None specified |
| Traffic statistics | <p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. | detail extensive |
| Protocol | Protocol family configured on the logical interface (such as iso or inet6). | detail extensive none |
| MTU | MTU size on the logical interface. | detail extensive none |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Addresses, Flags | Addresses associated with the logical interface and information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |

Table 123: Dynamic Flow Capture show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------|--|-----------------------|
| Destination | IP address of the remote side of the connection. | detail extensive none |
| Local | IP address of the logical interface. | detail extensive none |

Sample Output

show interfaces (Dynamic Flow Capture)

```

user@host> show interfaces dfc-0/0/0
Physical interface: dfc-0/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 36
  Type: Adaptive-Services, Link-level type: Dynamic-Flow-Capture, MTU: 9192, Speed:
  2488320kbps
  Device flags : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link type : Full-Duplex
  Link flags : None
  Last flapped : 2005-08-26 15:08:36 PDT (01:18:42 ago)
  Input rate : 0 bps (0 pps)
  Output rate : 44800440 bps (100000 pps)

Logical interface dfc-0/0/0.0 (Index 67) (SNMP ifIndex 43)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Dynamic-Flow-Capture
  Input packets : 74
  Output packets: 132
  Protocol inet, MTU: 9192
    Flags: Receive-options, Receive-TTL-Exceeded
    Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.36.100.1, Local: 10.36.100.2

Logical interface dfc-0/0/0.1 (Index 68) (SNMP ifIndex 49)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Dynamic-Flow-Capture
  Input packets : 0
  Output packets: 402927263
  Protocol inet, MTU: 9192
    Flags: Receive-options, Receive-TTL-Exceeded

Logical interface dfc-0/0/0.2 (Index 69) (SNMP ifIndex 50)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Dynamic-Flow-Capture
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 9192
    Flags: Receive-options, Receive-TTL-Exceeded

Logical interface dfc-0/0/0.16383 (Index 70) (SNMP ifIndex 44)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Dynamic-Flow-Capture
  Input packets : 1427
  Output packets: 98
  Protocol inet, MTU: 9192
    Flags: Receive-options, Receive-TTL-Exceeded
    Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.0.0.16, Local: 10.0.0.1

```


show interfaces (Flow Collector)

| | |
|---------------------------------|--|
| Syntax | <pre>show interfaces <i>cp-fpc/pic/port:channel</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M Series and T Series routers only) Display status information about the specified flow collector interface. |
| Options | <p><i>cp-fpc/pic/port:channel</i>—Display standard status information about the specified flow collector interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| List of Sample Output | show interfaces extensive (Flow Collector) on page 2187 |
| Output Fields | Table 124 on page 2183 lists the output fields for the show interfaces (Flow Collector) command. Output fields are listed in the approximate order in which they appear. |

Table 124: Flow Collector Show interfaces Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|---|------------------------------|
| Physical Interface | | |
| Physical Interface | Name of the physical interface type. | All levels |
| Link | Status of the link: up or down . | All levels |
| Enabled | State of the interface type. Possible values are described in the “Enabled Devices” section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |

Table 124: Flow Collector Show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|------------------------------|
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Type | Type of interface. | All levels |
| Link-level type | Encapsulation type used on the physical interface. | All levels |
| MTU | Maximum Transmit Unit (MTU). Size of the largest packet to be transmitted. | All levels |
| Clocking | Reference clock source of the interface. | All levels |
| Speed | Network speed on the interface. | All levels |
| Device flags | Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Link type | Data transmission type. | All levels |
| Link flags | Information about the link. Possible values are described in the "Link Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Physical info | Information about the physical interface. | All levels |
| Hold-times | Current interface hold-time up and hold-time down. Value is in milliseconds. | detail extensive none |
| Current address | Configured MAC address. | detail extensive none |
| Hardware address | Media access control (MAC) address of the interface. | detail extensive none |
| Alternate link address | Backup link address. | detail extensive none |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) . | detail extensive |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Traffic statistics | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. | detail extensive |

Table 124: Flow Collector Show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|--|------------------------------|
| Input errors | <ul style="list-style-type: none"> • Errors—Input errors on the interface. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Frames received smaller than the runt threshold. • Giants—Frames received larger than the giant threshold. • Policed Discards—Frames that the incoming packet match code discarded because the frames did not recognize them or were not of interest. Usually, this field reports protocols that Junos does not support. • Resource errors—Sum of transmit drops. | extensive |
| Output errors | <ul style="list-style-type: none"> • Carrier transitions —Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • Errors—Sum of outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet dropped by the ASIC RED mechanism. • Resource errors—Sum of transmit drops. | extensive |
| Logical Interface | | |
| Logical interface | Name of the logical interface | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Flags | Information about the logical interface; values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Encapsulation | Encapsulation on the logical interface. | All levels |
| Traffic statistics | <p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. | detail extensive |

Table 124: Flow Collector Show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|--|------------------------------|
| Local statistics | Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Transit statistics | Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Protocol | Protocol family configured on the logical interface (such as iso or inet6). | detail extensive none |
| MTU | MTU size on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route table | Route table in which this address exists; for example, Route table:0 refers to inet.0. | detail extensive |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Addresses, Flags | Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Destination | IP address of the remote side of the connection. | detail extensive none |
| Local | IP address of the logical interface. | detail extensive none |
| Broadcast | Broadcast address. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |

Sample Output

show interfaces extensive (Flow Collector)

```

user@host> show interfaces extensive cp-5/0/0
Physical interface: cp-5/0/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 52, Generation: 29
  Type: Flow-collector, Link-level type: Flow-collection, MTU: 9192,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2005-05-24 16:48:11 PDT (00:12:04 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :          2041661287          0 bps
    Output bytes :          3795049544      43816664 bps
    Input packets:          1365534          0 pps
    Output packets:          3865644      3670 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 2, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

Logical interface cp-5/0/0.0 (Index 74) (SNMP ifIndex 53) (Generation 28)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Flow-collection
  Traffic statistics:
    Input bytes :          1064651568
    Output bytes :           37144290
    Input packets:           711324
    Output packets:          713672
  Local statistics:
    Input bytes :              0
    Output bytes :              0
    Input packets:              0
    Output packets:              0
  Transit statistics:
    Input bytes :          1064651568          0 bps
    Output bytes :           37144290          0 bps
    Input packets:           711324          0 pps
    Output packets:          713672          0 pps
  Protocol inet, MTU: 9192, Generation: 39, Route table: 0
    Flags: Receive-options, Receive-TTL-Exceeded
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 4.0.0.2, Local: 4.0.0.1, Broadcast: Unspecified,
      Generation: 40

Logical interface cp-5/0/0.1 (Index 75) (SNMP ifIndex 54) (Generation 29)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Flow-collection
  Traffic statistics:
    Input bytes :          976793823
    Output bytes :          34099481
    Input packets:           652729
    Output packets:          655127

```

```
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 976793823 0 bps
Output bytes : 34099481 0 bps
Input packets: 652729 0 pps
Output packets: 655127 0 pps
Protocol inet, MTU: 9192, Generation: 40, Route table: 0
Flags: Receive-options, Receive-TTL-Exceeded
Addresses, Flags: Is-Preferred Is-Primary
Destination: 4.1.1.2, Local: 4.1.1.1, Broadcast: Unspecified,
Generation: 42

Logical interface cp-5/0/0.2 (Index 80) (SNMP ifIndex 55) (Generation 30)
Flags: Point-To-Point SNMP-Traps Encapsulation: Flow-collection
Traffic statistics:
Input bytes : 0
Output bytes : 3723079376
Input packets: 0
Output packets: 2495372
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 3723079376 43816664 bps
Input packets: 0 0 pps
Output packets: 2495372 3670 pps
Protocol inet, MTU: 9192, Generation: 41, Route table: 0
Flags: Receive-options, Receive-TTL-Exceeded
Addresses, Flags: Is-Preferred Is-Primary
Destination: 4.2.2.2, Local: 4.2.2.1, Broadcast: Unspecified,
Generation: 44

Logical interface cp-5/0/0.16383 (Index 81) (SNMP ifIndex 56) (Generation 31)
...
```

show interfaces (Flow Monitoring)

| | |
|---------------------------------|---|
| Syntax | <pre>show interfaces mo-fpc/pic/port:channel <brief detail extensive terse> <descriptions> <media> <snmp-index snmp-index> <statistics></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M Series and T Series routers only) Display status information about the specified flow monitoring interface. |
| Options | <p>mo-fpc/pic/port:channel—Display standard status information about the specified flow monitoring interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index snmp-index—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| List of Sample Output | show interfaces extensive (Flow Monitoring) on page 2192 |
| Output Fields | Table 125 on page 2189 lists the output fields for the show interfaces (Flow Monitoring) command. Output fields are listed in the approximate order in which they appear. |

Table 125: show interfaces Output Fields (Flow Monitoring)

| Field Name | Field Description | Level of Output |
|---------------------------|--|------------------------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |
| Link | Status of the link: up or down . | All levels |
| Enabled | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |

Table 125: show interfaces Output Fields (Flow Monitoring) (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|------------------------------|
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Description | Description and name of the interface. | All levels |
| Type | Type of interface. | All levels |
| Link-level type | Encapsulation type used on the physical interface. | All levels |
| MTU | Maximum Transmit Unit (MTU). Size of the largest packet to be transmitted. | All levels |
| Clocking | Reference clock source of the interface. | All levels |
| Speed | Network speed on the interface. | All levels |
| Device flags | Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Link type | Data transmission type. | All levels |
| Link flags | Information about the link. Possible values are described in the "Link Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Physical info | Information about the physical interface. | All levels |
| Hold-times | Current interface hold-time up and hold-time down. Value is in milliseconds. | detail extensive |
| Current address | Configured MAC address. | detail extensive none |
| Hardware address | Media access control (MAC) address of the interface. | detail extensive none |
| Alternate link address | Backup link address. | detail extensive none |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour:minute:second timezone (hour:minute:second ago)</i> . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) | detail extensive |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |

Table 125: show interfaces Output Fields (Flow Monitoring) (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|---|------------------------------|
| Traffic statistics | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. | detail extensive |
| Input errors | <ul style="list-style-type: none"> • Errors—Input errors on the interface. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Frames received smaller than the runt threshold. • Giants—Frames received larger than the giant threshold. • Policed Discards—Frames that the incoming packet match code discarded because the frames did not recognize them or were not of interest. Usually, this field reports protocols that Junos does not support. • Resource errors—Sum of transmit drops. | extensive |
| Output errors | <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • Errors—Sum of outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet dropped by the ASIC Red mechanism. • Resource errors—Sum of transmit drops. | extensive |
| Logical Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Flags | Information about the logical interface; values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Encapsulation | Encapsulation on the logical interface. | All levels |

Table 125: show interfaces Output Fields (Flow Monitoring) (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|---|------------------------------|
| Traffic statistics | <p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. | detail extensive |
| Local statistics | Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Transit statistics | Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Protocol | Protocol family configured on the logical interface (such as iso or inet6). | detail extensive none |
| MTU | MTU size on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route table | Route table in which this address exists; for example, Route table:0 refers to inet.0 . | detail extensive |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |

Sample Output

show interfaces extensive (Flow Monitoring)

```

user@host> show interfaces mo-4/0/0 extensive
Physical interface: mo-4/0/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 42, Generation: 28
  Description: monitor pic 2
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: Unlimited,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times    : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped  : 2005-05-24 16:43:12 PDT (00:17:46 ago)
  Statistics last cleared: Never

```

```
Traffic statistics:
Input bytes :          756824218          8328536 bps
Output bytes :          872916185          8400160 bps
Input packets:           508452           697 pps
Output packets:        15577196          18750 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
Policed discards: 0, Resource errors: 0
Output errors:
Carrier transitions: 2, Errors: 0, Drops: 0, MTU errors: 0,
Resource errors: 0

Logical interface mo-4/0/0.0 (Index 83) (SNMP ifIndex 43) (Generation 26)
Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
Traffic statistics:
Input bytes :          756781796
Output bytes :          872255328
Input packets:           507233
Output packets:        15575988
Local statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:             0
Transit statistics:
Input bytes :          756781796          8328536 bps
Output bytes :          872255328          8400160 bps
Input packets:           507233           697 pps
Output packets:        15575988          18750 pps
Protocol inet, MTU: Unlimited, Generation: 38, Route table: 0
Flags: None

Logical interface mo-4/0/0.16383 (Index 84) (SNMP ifIndex 58) (Generation 27)
...
```

show passive-monitoring error

| | |
|---------------------------------|--|
| Syntax | <code>show passive-monitoring error (* all mo-fpc/pic/port)</code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M40e, M160, and M320 routers and T Series routers only) Display passive monitoring error statistics. |
| Options | <code>* all mo-fpc/pic/port</code> —Display error statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name. |
| Required Privilege Level | view |
| List of Sample Output | show passive-monitoring error all on page 2195 |
| Output Fields | Table 126 on page 2194 lists the output fields for the show passive-monitoring error command. Output fields are listed in the approximate order in which they appear. |

Table 126: show passive-monitoring error Output Fields

| Field Name | Field Description |
|------------------------------------|---|
| Passive monitoring interface | Name of the passive monitoring interface. |
| Local interface index | Index counter of the local interface. |
| Interface state | State of the passive monitoring interface: <ul style="list-style-type: none"> • Monitoring—Specified interface is actively monitoring. • Disabled—Specified interface has been disabled from the CLI. • Not monitoring—The interface is operational, but not monitoring. This condition occurs when an interface first comes online, or when the interface is operational, but no logical unit has been configured under the physical interface. • Unknown—Unknown state. • Error—An error occurred during the process of determining the state of the interface. |
| Error information | |
| Packets dropped (no memory) | Number of packets dropped because of memory shortage. |
| Packets dropped (not IP) | Number of non-IP packets dropped. |
| Packets dropped (not IPv4) | Number of packets dropped because they failed the IPv4 version check. |
| Packets dropped (header too small) | Number of packets dropped because the packet length or IP header length was too small. |

Table 126: show passive-monitoring error Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------------------|---|
| Memory allocation failures | Number of flow record memory allocation failures. A small number reflects failures to replenish the free list. A large number indicates the monitoring station is almost out of memory space. |
| Memory free failures | Number of flow record memory free failures. |
| Memory free list failures | Number of flow records received from free list that failed. Memory is nearly exhausted or too many new flows greater than 128 KB are being created per second. |
| Memory warning | Whether the flows have exceeded 1 million packets per second (Mpps) on a Monitoring Services PIC or 2 Mpps on a Monitoring Services II PIC. The response can be Yes or No . |
| Memory overload | Whether the memory has been overloaded. The response can be Yes or No . |
| PPS overload | Whether the PIC is receiving more packets per second than the configured threshold. The response can be Yes or No . |
| BPS overload | Whether the PIC is receiving more bits per second than the configured threshold. The response can be Yes or No . |

Sample Output

show passive-monitoring error all

```

user@host> show passive-monitoring error all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Interface state: Monitoring
Error information
  Packets dropped (no memory): 0, Packets dropped (not IP): 0
  Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
  Memory allocation failures: 0, Memory free failures: 0
  Memory free list failures: 0
  Memory warning: No, Memory overload: No, PPS overload: No, BPS overload: No

Passive monitoring interface: mo-4/1/0, Local interface index: 45
Interface state: Not monitoring
Error information
  Packets dropped (no memory): 0, Packets dropped (not IP): 0
  Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
  Memory allocation failures: 0, Memory free failures: 0
  Memory free list failures: 0
  Memory warning: No, Memory overload: No, PPS overload: No, BPS overload: No

```

show passive-monitoring flow

| | |
|---------------------------------|---|
| Syntax | <code>show passive-monitoring flow (* all mo-<i>fpc/pic/port</i>)</code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M40e, M160, and M320 routers and T Series routers only) Display passive flow statistics. |
| Options | <code>* all mo-<i>fpc/pic/port</i></code> —Display passive flow statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name. |
| Required Privilege Level | view |
| List of Sample Output | show passive-monitoring flow all on page 2197 |
| Output Fields | Table 127 on page 2196 lists the output fields for the <code>show passive-monitoring flow</code> command. Output fields are listed in the approximate order in which they appear. |

Table 127: show passive-monitoring flow Output Fields

| Field Name | Field Description |
|------------------------------|---|
| Passive monitoring interface | Name of the passive monitoring interface. |
| Local interface index | Index counter of the local interface. |
| Interface state | State of the passive monitoring interface: <ul style="list-style-type: none"> • Monitoring—Specified interface is actively monitoring. • Disabled—Specified interface has been disabled from the CLI. • Not monitoring—The interface is operational, but not monitoring. This condition occurs when an interface first comes online, or when the interface is operational, but no logical unit has been configured under the physical interface. • Unknown—Unknown state. • Error—An error occurred during the process of determining the state of the interface. |
| Flow information | |
| Flow packets | Number of packets received by an operational PIC. |
| Flow bytes | Number of bytes received by an operational PIC. |
| Flow packets 10-second rate | Number of packets per second handled by the PIC and displayed as a 10-second average. |
| Flow bytes 10-second rate | Number of bytes per second handled by the PIC and displayed as a 10-second average. |
| Active flows | Number of currently active flows tracked by the PIC. |
| Total flows | Total number of flows received by an operational PIC. |

Table 127: show passive-monitoring flow Output Fields (*continued*)

| Field Name | Field Description |
|---------------------------------|--|
| Flows exported | Total number of flows exported by an operational PIC. |
| Flows packets exported | Total number of cflowd packets exported by an operational PIC. |
| Flows inactive timed out | Total number of flows that are exported because of inactivity. |
| Flows active timed out | Total number of long-lived flows that are exported because of an active timeout. |

Sample Output

show passive-monitoring flow all

```

user@host> show passive-monitoring flow all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Interface state: Monitoring
Flow information
  Flow packets: 6533434, Flow bytes: 653343400
  Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
  Active flows: 0, Total flows: 1599
  Flows exported: 1599, Flows packets exported: 55
  Flows inactive timed out: 1599, Flows active timed out: 0

Passive monitoring interface: mo-4/1/0, Local interface index: 45
Interface state: Monitoring
Flow information
  Flow packets: 6537780, Flow bytes: 653778000
  Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
  Active flows: 0, Total flows: 1601
  Flows exported: 1601, Flows packets exported: 55
  Flows inactive timed out: 1601, Flows active timed out: 0

```

show passive-monitoring memory

| | |
|---------------------------------|---|
| Syntax | <code>show passive-monitoring memory (* all mo-<i>fpc/pic/port</i>)</code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M40e, M160, and M320 routers and T Series routers only) Display passive monitoring memory and flow record statistics |
| Options | <code>* all mo-<i>fpc/pic/port</i></code> —Display memory and flow record statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name. |
| Required Privilege Level | view |
| List of Sample Output | show passive-monitoring memory all on page 2198 |
| Output Fields | Table 128 on page 2198 lists the output fields for the show passive-monitoring memory command. Output fields are listed in the approximate order in which they appear. |

Table 128: show passive-monitoring memory Output Fields

| Field Name | Field Description |
|---|---|
| Passive monitoring interface | Name of the passive monitoring interface. |
| Local interface index | Index counter of the local interface. |
| Memory utilization | |
| Allocation count | Number of flow records allocated. |
| Free count | Number of flow records freed. |
| Maximum allocated | Maximum number of flow records allocated since the monitoring station booted. This number represents the peak number of flow records allocated at a time. |
| Allocations per second | Flow records allocated per second during the last statistics interval on the PIC. |
| Frees per second | Flow records freed per second during the last statistics interval on the PIC. |
| Total memory used,
Total memory free | Total memory currently used and total amount of memory currently free (in bytes). |

Sample Output

show passive-monitoring memory all

```
user@host> show passive-monitoring memory all
```



```
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Memory utilization
  Allocation count: 1600, Free count: 1599, Maximum allocated: 1600
  Allocations per second: 3200, Frees per second: 1438
  Total memory used (in bytes): 103579176, Total memory free (in bytes):
  163914184
```

show passive-monitoring status

| | |
|---------------------------------|---|
| Syntax | <code>show passive-monitoring status (* all mo-fpc/pic/port)</code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M40e, M160, and M320 routers and T Series routers only) Display passive monitoring status. |
| Options | <code>* all mo-fpc/pic/port</code> —Display status for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name. |
| Required Privilege Level | view |
| List of Sample Output | show passive-monitoring status all on page 2201 |
| Output Fields | Table 129 on page 2200 lists the output fields for the show passive-monitoring status command. Output fields are listed in the approximate order in which they appear. |

Table 129: show passive-monitoring status Output Fields

| Output Field | Output Field Description |
|------------------------------|--|
| Passive monitoring interface | Name of the passive monitoring interface. |
| Local interface index | Index counter of the local interface. |
| Interface state | Monitoring state of the passive monitoring interface. <ul style="list-style-type: none"> • Monitoring—PIC is actively monitoring. • Disabled—PIC has been disabled using the CLI. • Not monitoring—PIC is operational, but not monitoring. This condition can happen while the PIC is coming online, or when the PIC is operational but has no logical unit configured under the physical interface. • Unknown |
| Group index | Integer that represents the monitoring group of which the PIC is a member. Group index is a mapping from the group name to an index. It is not related to the number of monitoring groups. |
| Export interval | Configured export interval for cflowd records, in seconds. |
| Export format | Configured export format (only cflowd version 5 is supported). |
| Protocol | Protocol the PIC is configured to monitor (only IPv4 is supported). |
| Engine type | Configured engine type that is inserted in output cflowd packets. |
| Engine ID | Configured engine ID that is inserted in output cflowd packets. |

Sample Output

show passive-monitoring status all

```
user@host> show passive-monitoring status all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Interface state: Monitoring
  Group index: 0
  Export interval: 15 secs, Export format: cflowd v5
  Protocol: IPv4, Engine type: 1, Engine ID: 1

Passive monitoring interface: mo-4/1/0, Local interface index: 45
Interface state: Disabled

Passive monitoring interface: mo-4/2/0, Local interface index: 46
Interface state: Not monitoring
```

show passive-monitoring usage

| | |
|---------------------------------|--|
| Syntax | <code>show passive-monitoring usage (* all mo-fpc/pic/port)</code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M40e, M160, and M320 routers and T Series routers only) Display passive monitoring usage statistics. |
| Options | <code>* all mo-fpc/pic/port</code> —Display usage statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name. |
| Required Privilege Level | view |
| List of Sample Output | show passive-monitoring usage all on page 2202 |
| Output Fields | Table 130 on page 2202 lists the output fields for the show passive-monitoring usage command. Output fields are listed in the approximate order in which they appear. |

Table 130: show passive-monitoring usage Output Fields

| Output Field | Output Field Description |
|------------------------------|--|
| Passive monitoring interface | Name of the passive monitoring interface. |
| Local interface index | Index counter of the local interface. |
| CPU utilization | |
| Uptime | Time, in milliseconds, that the PIC has been operational. |
| Interrupt time | Total time that the PIC has spent processing packets since the last PIC reset. |
| Load (5 second) | CPU load on the PIC, averaged more than 5 seconds. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time. |
| Load (1 minute) | CPU load on the PIC, averaged more than 1 minute. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time. |

Sample Output

show passive-monitoring usage all

```

user@host> show passive-monitoring usage
Passive monitoring interface: mo-4/0/0, Local interface index: 44
CPU utilization
  Uptime: 653155 milliseconds, Interrupt time: 40213754 microseconds
  Load (5 second): 20%, Load (1 minute): 17%

Passive monitoring interface: mo-4/1/0, Local interface index: 45
CPU utilization

```

Uptime: 652292 milliseconds, Interrupt time: 40223178 microseconds
Load (5 second): 22%, Load (1 minute): 15%

Passive monitoring interface: mo-4/2/0, Local interface index: 46

CPU utilization

Uptime: 649491 milliseconds, Interrupt time: 40173645 microseconds
Load (5 second): 22%, Load (1 minute): 10098862%

show services accounting aggregation

| | |
|---------------------------------|---|
| Syntax | <code>show services accounting aggregation <i>aggregation-type</i> <<i>aggregation-value</i>></code>
<code><detail extensive terse></code>
<code><limit <i>limit-value</i>></code>
<code>< name <i>service-name</i>></code>
<code><order (bytes packets)></code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display information about the aggregated active flows being processed by the accounting service. |
| Options | <p><i>aggregation-type</i> <<i>aggregation-value</i>>—Display information for a particular aggregation type and optional value:</p> <ul style="list-style-type: none"><i>as</i> <<i>source-as-value</i> <i>destination-as-value</i> <i>input-snmp-interface-index-value</i> <i>output-snmp-interface-index-value</i>>—Aggregate by autonomous system (AS).<i>destination-prefix</i> <<i>destination-prefix-value</i> <i>destination-as-value</i> <i>output-snmp-interface-index-value</i>>—Aggregate by destination prefix.<i>protocol-port</i> <<i>protocol-value</i> <i>source-port-value</i> <i>destination-port-value</i>>—Aggregate by protocol and port.<i>source-destination-prefix</i> <<i>source-prefix-value</i> <i>destination-prefix-value</i> <i>destination-as-value</i> <i>source-as-value</i> <i>input-snmp-interface-index-value</i> <i>output-snmp-interface-index-value</i>>—Aggregate by source and destination prefix.<i>source-prefix</i> <<i>source-prefix-value</i> <i>source-as-value</i> <i>input-snmp-interface-index-value</i>>—Aggregate by source prefix. <p>detail extensive terse—(Optional) Display the specified level of output.</p> <p>limit <i>limit-value</i>—(Optional) Limit the display output to this number of flows. The default is no limit.</p> <p>name <i>service-name</i>—(Optional) Display information about the aggregated flows for a particular service name.</p> <p>order (bytes packets)—(Optional) Display the flow with the ordering of the highest number, either by byte count or by packet count.</p> |
| Additional Information | For information about aggregation configuration options, see the <i>Junos OS Services Interfaces Library for Routing Devices</i> . |
| Required Privilege Level | view |
| List of Sample Output | show services accounting aggregation protocol-port detail on page 2206
show services accounting aggregation source-destination-prefix on page 2206 |

[show services accounting aggregation source-destination- prefix order packet detail on page 2206](#)

[show services accounting aggregation source-destination- prefix extensive limit on page 2207](#)

[show services accounting aggregation source-destination-prefix name terse on page 2207](#)

Output Fields [Table 131 on page 2205](#) lists the output fields for the **show services accounting aggregation** command. Output fields are listed in the approximate order in which they appear.

Table 131: show services accounting aggregation Output Fields

| Field Name | Field Description |
|------------------------------|---|
| Service Accounting interface | Name of the service accounting interface. |
| Local interface index | Index corresponding to the service accounting interface. |
| Service name | Name of a service that was configured at the [edit forwarding-options accounting] hierarchy level. The default display, (default sampling), indicates the service was configured at the [edit forwarding-options sampling-level] hierarchy level. |
| Protocol | Protocol identifier and number. |
| Source Port | Source port identifier and number. |
| Destination Port | Destination port identifier and number. |
| Source-AS | Source autonomous system (AS) number. |
| Destination-AS | Destination AS number. |
| Source Prefix | Source prefix. |
| Destination Prefix | Destination prefix. |
| Source address | Source address. |
| Source prefix length | Source prefix length. |
| Destination address | Destination address. |
| Destination prefix length | Destination prefix length. |
| Input SNMP interface index | SNMP index of the interface the packet came in on. |
| Output SNMP interface index | SNMP index of the interface the packet went out on. |

Table 131: show services accounting aggregation Output Fields (*continued*)

| Field Name | Field Description |
|--------------|---|
| Start time | Actual time when the packet in this aggregation was first seen. |
| End time | Actual time when the packet in this aggregation was last seen. |
| Flow count | Number of flows in the aggregation. |
| Packet count | Number of packets in the aggregation. |
| Byte count | Number of bytes in the aggregation. |

Sample Output

show services accounting aggregation protocol-port detail

```

user@host> show service accounting aggregation protocol-port detail
Service Accounting interface: mo-2/0/0, Local interface index: 468
Service name: (default sampling)
  Protocol: 6, Source port: 20, Destination port: 20
  Start time: 442349, End time: 6425714
  Flow count: 194, Packet count: 4294964388, Byte count: 4294781184

  Protocol: 0, Source port: 0, Destination port: 0
  Start time: 442349, End time: 6425749
  Flow count: 204, Packet count: 4294964324, Byte count: 4294777088

  Protocol: 17, Source port: 123, Destination port: 123
  Start time: 442364, End time: 6425784
  Flow count: 186, Packet count: 4294964152, Byte count: 4294766080

```

show services accounting aggregation source-destination-prefix

```

user@host> show service accounting aggregation source-destination-prefix
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
Source          Destination    Input          Output          Flow    Packet
Byte           prefix        interface      interface      count   count
prefix        count
11.1.0.0/20    40.0.0.0/24   ge-5/0/1.0     ge-5/0/0.0     256     491761
31472704
11.1.0.0/20    40.0.1.36/32  ge-5/0/1.0     ge-5/0/0.0     1
1926          123264
11.1.0.0/20    40.0.1.59/32  ge-5/0/1.0     ge-5/0/0.0     1
1926          123264
11.1.0.0/20    40.0.3.63/32  ge-5/0/1.0     ge-5/0/0.0     1
1925          123200
11.1.0.0/20    40.0.3.32/32  ge-5/0/1.0     ge-5/0/0.0     1
1925

```

show services accounting aggregation source-destination- prefix order packet detail

```

user@host> show service accounting aggregation source-destination-prefix order packet detail
name t2 input-snmp-interface-index 538

```



```

Service Accounting interface: mo-2/0/0, Local interface index: 468
Service name: t2
Source      Destination  Input SNMP  Output SNMP  Flow  Packet  Byte
Prefix      Prefix      Index      Index      Count Count   Count
11.1.1.2/20 30.0.167.1/0 538        432         1     60     46483
11.1.1.2/20 30.0.168.1/0 538        432         1     60     5191
11.1.1.2/20 30.0.154.1/0 538        432         2     60     45504
11.1.1.2/20 30.0.76.1/0  538        432         1     60     42177
11.1.1.2/20 30.0.149.1/0 538        432         1     60     49184
11.1.1.2/20 30.0.113.1/0 538        432         2     60     48757

```

show services accounting aggregation source-destination- prefix extensive limit

```

user@host> show service accounting aggregation source-destination-prefix name t2 extensive
limit 3

```

```

Service Accounting interface: mo-2/0/0, Local interface index: 542
Service name: t2

```

```

Source address: 11.1.1.2, Source prefix length: 20
Destination address: 44.200.176.1, Destination prefix length: 0
Input SNMP interface index: 24, Output SNMP interface index: 26
Source-AS: 69, Destination-AS: 69
Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003
Flow count: 0, Packet count: 6, Byte count: 5340

```

```

Source address: 11.1.1.2, Source prefix length: 20
Destination address: 45.243.160.1, Destination prefix length: 0
Input SNMP interface index: 24, Output SNMP interface index: 26
Source-AS: 69, Destination-AS: 69
Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003
Flow count: 0, Packet count: 6, Byte count: 5490

```

```

Source address: 11.1.1.2, Source prefix length: 20
Destination address: 45.162.160.1, Destination prefix length: 0
Input SNMP interface index: 24, Output SNMP interface index: 26
Source-AS: 69, Destination-AS: 69
Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003
Flow count: 0, Packet count: 6, Byte count: 4079

```

show services accounting aggregation source-destination-prefix name terse

```

user@host> show service accounting aggregation source-destination-prefix name T3 terse

```

```

Service Accounting interface: rsp0, Local interface index: 171

```

```

Service name: T3

```

```

Interface state: Accounting

```

| Source | Destination | Input | Output | Flow | Packet |
|-------------|--------------|------------|------------|-------|--------|
| Byte | | | | | |
| prefix | prefix | interface | interface | count | count |
| 11.1.0.0/20 | 50.0.0.0/24 | ge-5/0/1.0 | ge-5/0/0.0 | 256 | 639822 |
| 40948608 | | | | | |
| 11.1.0.0/20 | 50.0.2.67/32 | ge-5/0/1.0 | ge-5/0/0.0 | 1 | |
| 2485 | 159040 | | | | |
| 11.1.0.0/20 | 50.0.2.92/32 | ge-5/0/1.0 | ge-5/0/0.0 | 1 | |
| 2485 | | | | | |

show services accounting aggregation template

| | |
|---------------------------------|--|
| Syntax | show services accounting aggregation template
<template-name <i>template-name</i>> |
| Release Information | Command introduced in Junos OS Release 8.3. |
| Description | Display information for flow aggregation version 9 templates. |
| Options | <template-name <i>template-name</i>> —(Optional) Display information for the specified template only. |
| Required Privilege Level | view |
| List of Sample Output | show services accounting aggregation template on page 2208 |
| Output Fields | Table 132 on page 2208 lists the output fields for the show services accounting aggregation template command. Output fields are listed in the approximate order in which they appear. |

Table 132: show services accounting aggregation template Output Fields

| Field Name | Field Description |
|-------------------------------|---------------------------------|
| MPLS Label 1 | Position of first MPLS label. |
| MPLS Label 2 | Position of second MPLS label. |
| MPLS Label 3 | Position of third MPLS label. |
| MPLS Top Level Address | Outer top label FEC IP address. |
| Packet Count | Number of packets sent. |

Sample Output

show services accounting aggregation template

```

user@host> show services accounting aggregation template template-name mpls
MPLS label 1: 299808, MPLS label 2: 0, MPLS label 3: 0
Source address: 11.1.1.2, Destination address: 10.255.15.22, Top Label Address:
22.15.255.10
Source port: 0, Destination port: 0
Protocol: 61, TOS: 0, TCP flags: 0
Source mask: 24, Destination mask: 32
Input SNMP interface index: 503, Output SNMP interface index: 505
Start time: 40780, End time: 157330
Packet count: 3949198, Byte count: 181663062

```

show services accounting errors

| | |
|---------------------------------|--|
| Syntax | show services accounting errors
<inline-jflow name (* all <i>service-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display active flow error statistics. |
| Options | <p>none—Display error statistics for all services accounting instances.</p> <p>inline-jflow fpc-slot <i>slot-number</i>—(Optional) Display error statistics for inline jflow.</p> <p>name (* all <i>service-name</i>)—(Optional) Display active flow error statistics. Use a wildcard character, specify all services, or provide a specific service name.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> show services accounting flow on page 2213 |
| List of Sample Output | <p>show services accounting errors (Monitoring PIC interface) on page 2210</p> <p>show services accounting errors (Service PIC interface) on page 2211</p> <p>show services accounting errors inline-jflow fpc-slot slot-number (when only IPv6 is configured) on page 2211</p> <p>show services accounting errors inline-jflow fpc-slot slot-number (when both IPv4 and IPv6 are configured) on page 2211</p> <p>show services accounting errors inline-jflow (MX80 Router when both IPv4 and IPv6 are configured) on page 2211</p> |
| Output Fields | Table 133 on page 2209 lists the output fields for the show services accounting errors command. Output fields are listed in the approximate order in which they appear. |

Table 133: show services accounting errors Output Fields

| Field | Field Description |
|------------------------------|---|
| Service Accounting interface | Name of the service accounting interface. |
| Local interface index | Index counter of the local interface. |
| FPC slot | Slot number of the FPC for which the flow information is displayed. (Available only when the inline-jflow fpc-slot slot-number option is used.) |
| Service name | Name of a service that was configured at the [edit forwarding-options accounting] hierarchy level. The default display, (default sampling) , indicates the service was configured at the [edit forwarding-options sampling-level] hierarchy level. |

Error Information

Table 133: show services accounting errors Output Fields (*continued*)

| Field | Field Description |
|------------------------------------|---|
| Packets dropped (no memory) | Number of packets dropped because of memory shortage. |
| Packets dropped (not IP) | Number of non-IP packets dropped. |
| Packets dropped (not IPv4) | Number of packets dropped because they failed the IPv4 version check. |
| Packets dropped (header too small) | Number of packets dropped because the packet length or IP header length was too small. |
| Memory allocation failures | Number of flow record memory allocation failures. A small number reflects failures to replenish the free list. A large number indicates the monitoring station is almost out of memory space. |
| Memory free failures | Number of flow record memory free failures. |
| Memory free list failures | Number of flow records received from the free list that failed. Memory is nearly exhausted, or too many new flows greater than 128 KB are being created per second. |
| Memory overload | Whether the memory has been overloaded. The response can be Yes or No . |
| PPS overload | Whether the PIC is receiving more packets per second than the configured threshold. The response can be Yes or No . |
| BPS overload | Whether the PIC is receiving more bits per second than the configured threshold. The response can be Yes or No . |
| Flow Creation Failures | Number of times flow creation failed. |
| Route Record Lookup Failures | Number of times the route record lookup failed. |
| AS Lookup Failures | Number of times autonomous system lookup failed. |
| Export Packet Failures | Number of times packet export failed. |

Sample Output

show services accounting errors (Monitoring PIC interface)

```

user@host> show services accounting errors
Service Accounting interface: mo-1/1/0, Local interface index: 15
Service name: (default sampling)
Error information
  Packets dropped (no memory): 0, Packets dropped (not IP): 0
  Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
  Memory allocation failures: 0, Memory free failures: 0
  Memory free list failures: 0
  Memory overload: No, PPS overload: No, BPS overload: No

```

Sample Output

show services accounting errors (Service PIC interface)

```

user@host> show services accounting errors
Service Accounting interface: sp-0/1/0
Service name: (default sampling)
Error information
  Service sets dropped: 0, Active timeout failures: 0
  Export packet failures: 0, Flow creation failures: 0
  Memory overload: No

Service Accounting interface: sp-1/0/0
Service name: (default sampling)
Error information
  Service sets dropped: 0, Active timeout failures: 0
  Export packet failures: 0, Flow creation failures: 0
  Memory overload: No

```

show services accounting errors inline-jflow fpc-slot slot-number (when only IPv6 is configured)

```

user@host> show services accounting errors inline-jflow fpc-slot 5
Error information
  FPC Slot: 5
  Flow Creation Failures: 0
  Route Record Lookup Failures: 0, AS Lookup Failures: 0
  Export Packet Failures: 0
  Memory Overload: No, Memory Alloc Fail Count: 0

```

show services accounting errors inline-jflow fpc-slot slot-number (when both IPv4 and IPv6 are configured)

```

user@host> show services accounting errors inline-jflow fpc-slot 5
Error information
  FPC Slot: 5
  Flow Creation Failures: 0
  Route Record Lookup Failures: 0, AS Lookup Failures: 0
  Export Packet Failures: 0
  Memory Overload: No, Memory Alloc Fail Count: 0

IPv4:
IPv4 Flow Creation Failures: 0
IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
IPv4 Export Packet Failures: 0

IPv6:
IPv6 Flow Creation Failures: 0
IPv6 Route Record Lookup Failures: 0, IPv6 AS Lookup Failures: 0
IPv6 Export Packet Failures: 0

```

show services accounting errors inline-jflow (MX80 Router when both IPv4 and IPv6 are configured)

```

user@host> show services accounting errors inline-jflow
Error information
  TFEB Slot: 0
  Flow Creation Failures: 0
  Route Record Lookup Failures: 0, AS Lookup Failures: 0
  Export Packet Failures: 0
  Memory Overload: No

IPv4:
IPv4 Flow Creation Failures: 0

```

IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
IPv4 Export Packet Failures: 0

IPv6:

IPv6 Flow Creation Failures: 0
IPv6 Route Record Lookup Failures: 0, IPv6 AS Lookup Failures: 0
IPv6 Export Packet Failures: 0

show services accounting flow

| | |
|---------------------------------|--|
| Syntax | <code>show services accounting flow</code>
<code><inline-jflow logical-system name (* all service-name)></code> |
| Release Information | Command introduced before Junos OS Release 7.4.
Junos OS Release 10.0 added the capability to display output from multiple sampling instances. |
| Description | Display active flow statistics. |
| Options | <p>none—Display active flow statistics for all service instances.</p> <p>logical-system (all logical-system)—(Optional) Display active flow statistics for the specified logical system or all logical systems on the device.</p> <p>inline-jflow (fpc-slot slot-number)—(Optional) Display inline flow statistics for the specified FPC.</p> <p>name (* all service-name)—(Optional) Display services accounting active flow statistics. Use a wildcard character, specify all services, or provide a specific service name.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> show services accounting status on page 2227 |
| List of Sample Output | show services accounting flow (flow aggregation v5/v8 configuration) on page 2214
show services accounting flow (flow aggregation v9 configuration) on page 2214
show services accounting flow name on page 2215
show services accounting flow name all on page 2215
show services accounting flow (multiple sampling instances) on page 2216
show services accounting flow inline-jflow fpc-slot slot-number (for IPv4 flow) on page 2216
show services accounting flow inline-jflow fpc-slot slot-number (with IPv4 and IPv6 Configuration) on page 2216
show services accounting flow inline-jflow (MX80 Router with IPv4 and IPv6 Configuration) on page 2216 |
| Output Fields | Table 134 on page 2213 lists the output fields for the show services accounting flow command. Output fields are listed in the approximate order in which they appear. |

Table 134: show services accounting flow Output Fields

| Output Field | Output Field Description |
|------------------------------|---|
| Service Accounting interface | Name of the service accounting interface. |
| Local interface index | Index counter of the local interface. |

Table 134: show services accounting flow Output Fields (*continued*)

| Output Field | Output Field Description |
|------------------------------------|--|
| Service name | Name of a service that was configured at the [edit forwarding-options accounting] hierarchy level. The default display, (default sampling), indicates the service was configured at the [edit forwarding-options sampling-level] hierarchy level. |
| Flow Information | |
| FPC Slot | Slot number of the FPC for which the flow information is displayed. (Available only when the inline-jflow fpc-slot slot-number option is used.) |
| Flow packets | Number of packets received by an operational PIC. |
| Flow bytes | Number of bytes received by an operational PIC. |
| Flow packets 10-second rate | Number of packets per second handled by the PIC and displayed as a 10-second average. |
| Flow bytes 10-second rate | Number of bytes per second handled by the PIC and displayed as a 10-second average. |
| Active flows | Number of currently active flows tracked by the PIC. |
| Total flows | Total number of flows received by an operational PIC. |
| Flows exported | Total number of flows exported by an operational PIC. |
| Flows packets exported | Total number of cflowd packets exported by an operational PIC. |
| Flows inactive timed out | Total number of flows that are exported because of inactivity. |
| Flows active timed out | Total number of long-lived flows that are exported because of an active timeout. |

Sample Output

show services accounting flow (flow aggregation v5/v8 configuration)

```

user@host> show services accounting flow
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
Flow information
  Flow packets: 87168293, Flow bytes: 5578770752
  Flow packets 10-second rate: 45762, Flow bytes 10-second rate: 2928962
  Active flows: 1000, Total flows: 2000
  Flows exported: 19960, Flows packets exported: 582
  Flows inactive timed out: 1000, Flows active timed out: 29000

```

show services accounting flow (flow aggregation v9 configuration)

```

user@host> show services accounting flow
Flow information
  Service Accounting interface: sp-7/1/0, Local interface index: 149

```



```

Flow packets: 0, Flow bytes: 0
Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
Active flows: 0, Total flows: 0
Flows exported: 0, Flows packets exported: 1
Flows inactive timed out: 0, Flows active timed out: 0

```

show services accounting flow name

```

user@host> show services accounting flow count2
Service Accounting interface: mo-1/1/0, Local interface index: 15
Service name: count2
Flow information
  Flow packets: 0, Flow bytes: 0
  Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
  Active flows: 0, Total flows: 0
  Flows exported: 0, Flows packets exported: 0
  Flows inactive timed out: 0, Flows active timed out: 0

```

show services accounting flow name all

```

user@host> show services accounting flow name all
Service Accounting interface: rsp0, Local interface index: 171
Service name: T2
Interface state: Accounting
Flow information
  Flow packets: 37609891, Flow bytes: 2407033024
  Flow packets 10-second rate: 45762, Flow bytes 10-second rate: 2928953
  Active flows: 1000, Total flows: 1000
  Flows exported: 6705, Flows packets exported: 198
  Flows inactive timed out: 0, Flows active timed out: 13000

Service Accounting interface: rsp0, Local interface index: 171
Service name: T3
Interface state: Accounting
Flow information
  Flow packets: 37750807, Flow bytes: 2416051712
  Flow packets 10-second rate: 45762, Flow bytes 10-second rate: 2928940
  Active flows: 1000, Total flows: 1000
  Flows exported: 13437, Flows packets exported: 378
  Flows inactive timed out: 0, Flows active timed out: 13000

Service Accounting interface: rsp0, Local interface index: 171
Service name: T4
Interface state: Accounting
Flow information
  Flow packets: 0, Flow bytes: 0
  Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
  Active flows: 0, Total flows: 0
  Flows exported: 0, Flows packets exported: 0
  Flows inactive timed out: 0, Flows active timed out: 0

Service Accounting interface: rsp0, Local interface index: 171
Service name: count1
Interface state: Accounting
Flow information
  Flow packets: 0, Flow bytes: 0
  Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
  Active flows: 0, Total flows: 0
  Flows exported: 0, Flows packets exported: 0
  Flows inactive timed out: 0, Flows active timed out: 0

```

show services accounting flow (multiple sampling instances)

```
user@host> show services accounting flow
Flow information
Service Accounting interface: sp-2/0/0, Local interface index: 215
Flow packets: 9867, Flow bytes: 631488
Flow packets 10-second rate: 0, Flow bytes 10-second rate: 628
Active flows: 2, Total flows: 10
Flows exported: 4028, Flows packets exported: 6150
Flows inactive timed out: 8, Flows active timed out: 4026

Service Accounting interface: sp-2/1/0, Local interface index: 223
Flow packets: 0, Flow bytes: 0
Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
Active flows: 0, Total flows: 0
Flows exported: 0, Flows packets exported: 1
Flows inactive timed out: 0, Flows active timed out: 0
```

show services accounting flow inline-jflow fpc-slot slot-number (for IPv4 flow)

```
user@host> show services accounting flow inline-jflow fpc-slot 5
Flow information
FPC Slot: 5
Flow Packets: 0, Flow Bytes: 0
Active Flows: 0, Total Flows: 0
Flows Exported: 0, Flow Packets Exported: 0
Flows Inactive Timed Out: 0, Flows Active Timed Out: 0
```

show services accounting flow inline-jflow fpc-slot slot-number (with IPv4 and IPv6 Configuration)

```
user@host> show services accounting flow inline-jflow fpc-slot 5
Flow information
FPC Slot: 5
Flow Packets: 0, Flow Bytes: 0
Active Flows: 0, Total Flows: 0
Flows Exported: 0, Flow Packets Exported: 0
Flows Inactive Timed Out: 0, Flows Active Timed Out: 0

IPv4 Flows:
IPv4 Flow Packets: 0, IPv4 Flow Bytes: 0
IPv4 Active Flows: 0, IPv4 Total Flows: 0
IPv4 Flows Exported: 0, IPv4 Flow Packets exported: 0
IPv4 Flows Inactive Timed Out: 0, IPv4 Flows Active Timed Out: 0

IPv6 Flows:
IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
IPv6 Active Flows: 0, IPv6 Total Flows: 0
IPv6 Flows Exported: 0, IPv6 Flow Packets Exported: 0
IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0
```

show services accounting flow inline-jflow (MX80 Router with IPv4 and IPv6 Configuration)

```
user@host> show services accounting flow inline-jflow
Flow information
TFEB Slot: 0
Flow Packets: 0, Flow Bytes: 0
Active Flows: 0, Total Flows: 0
Flows Exported: 0, Flow Packets Exported: 0
Flows Inactive Timed Out: 0, Flows Active Timed Out: 0

IPv4 Flows:
```

IPv4 Flow Packets: 0, IPv4 Flow Bytes: 0
IPv4 Active Flows: 0, IPv4 Total Flows: 0
IPv4 Flows Exported: 0, IPv4 Flow Packets exported: 0
IPv4 Flows Inactive Timed Out: 0, IPv4 Flows Active Timed Out: 0

IPv6 Flows:
IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
IPv6 Active Flows: 0, IPv6 Total Flows: 0
IPv6 Flows Exported: 0, IPv6 Flow Packets Exported: 0
IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0

show services accounting flow-detail

Syntax `show services accounting flow-detail`
 `<detail | extensive | terse>`
 `<filters>`
 `<limit limit-value>`
 `<name (* | all | service-name)>`
 `<order (bytes | packets)>`

Release Information Command introduced before Junos OS Release 7.4.

Description Display information about the flows being processed by the accounting service.

Options `detail | extensive | terse`—(Optional) Display the specified level of output.

filters—(Optional) Filter the display output of the currently active flow records. The following filters query actively changing data structures and result in different results for multiple invocations:

- **destination-as**—Display flow records filtered by destination autonomous system information.
- **destination-port**—Display flow records filtered by destination port information.
- **destination-prefix**—Display flow records filtered by destination prefix information.
- **input-snmp-interface-index**—Display flow records filtered by SNMP input interface index information.
- **output-snmp-interface-index**—Display flow records filtered by SNMP output interface index information.
- **proto**—Display flow records filtered by protocol type.
- **source-as**—Display flow records filtered by source autonomous system information.
- **source-port**—Display flow records filtered by source port information.
- **source-prefix**—Display flow records filtered by source prefix information.
- **tos**—Display flow records filtered by type of service classification.

limit *limit-value*—(Optional) Limit the display output to the specified number of flows. The default is no limit.

name (* | all | *service-name*)—(Optional) Display information about the flows being processed. Use a wildcard character, specify all services, or provide a specific services name.

order (bytes | packets)—(Optional) Display the flow with the ordering of the highest number, either by byte count or by packet count.

Additional Information When no PIC is active, or when no route record has been downloaded from the PIC, this command reports no flows, even though packets are being sampled. This command

displays information about two concurrent sessions only. If a third session is attempted, the command pauses with no output until one of the previous sessions is completed.

Required Privilege Level view

List of Sample Output [show services accounting flow-detail on page 2220](#)
[show services accounting flow-detail limit on page 2221](#)
[show services accounting flow-detail name extensive on page 2221](#)
[show services accounting flow-detail limit order bytes on page 2221](#)
[show services accounting flow-detail source-port on page 2222](#)

Output Fields [Table 135 on page 2219](#) lists the output fields for the **show services accounting flow-detail** command. Output fields are listed in the approximate order in which they appear.

Table 135: show services accounting flow-detail Output Fields

| Field Name | Field Description | Output Level |
|-------------------------------------|---|------------------|
| Service Accounting interface | Name of the service accounting interface. | All levels |
| Service name | Name of a service that was configured at the [edit forwarding-options accounting] hierarchy level. The default display, (default sampling) , indicates the service was configured at the [edit forwarding-options sampling] hierarchy level. | All levels |
| Local interface index | Index counter of the local interface. | All levels |
| TOS | Type-of-service value from the IP header. | extensive |
| Input SNMP interface index | SNMP index of the interface on which the packet came in. | extensive |
| Output SNMP interface index | SNMP index of the interface on which the packet went out. | extensive |
| Source-AS | Source AS number. | extensive |
| Destination-AS | Destination AS number. | extensive |
| Protocol | Name of the protocol used for the packet flow from the corresponding source address. | All levels |
| Input interface | Interface on which the packets were received. | All levels |
| Output interface | Interface on which the packets were transmitted. | All levels |
| TCP flags | Number of TCP header flags detected in the flow. | extensive |
| Source address | Address where the flow originated. | All levels |
| Source port | Name of the source port. | All levels |

Table 135: show services accounting flow-detail Output Fields (*continued*)

| Field Name | Field Description | Output Level |
|--------------------------------------|---|------------------|
| Source prefix length | Source prefix length. | extensive |
| Destination address | Address where the flow is sent. | All levels |
| Destination prefix length | Destination prefix length. | extensive |
| Destination port | Name of the destination port. | All levels |
| Start time | Actual time when the packet in this aggregation was first seen. | detail extensive |
| End time | Actual time when the packet in this aggregation was last seen. | detail extensive |
| Packet count | Number of packets in the aggregation. | All levels |
| Byte count | Number of bytes in the aggregation. | All levels |
| Time since last active timeout | Amount of time elapsed since the last active timeout, in the format <i>hh:mm:ss</i> . | None specified |
| Packet count for last active timeout | Number of packets in the aggregation since the last active timeout. | None specified |
| Byte count for last active timeout | Number of bytes in the aggregation since the last active timeout. | None specified |

Sample Output

show services accounting flow-detail

In this sample, the output is split into three sections, with ellipses (...) indicating where the sections are continued.

```

user@host> show services accounting flow-detail
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting

```

| Protocol | Input interface | Source address | Source port | Output interface... |
|----------|-----------------|----------------|-------------|---------------------|
| tcp(6) | ge-5/0/1.0 | 11.1.1.2 | 0 | ge-5/0/0.0 |
| tcp(6) | ge-5/0/1.0 | 11.1.1.2 | 0 | ge-5/0/0.0 |

| Destination address | Destination port | Packet count | Byte count | Time since last active timeout... |
|---------------------|------------------|--------------|------------|-----------------------------------|
| 40.0.3.149 | 0 | 2660 | 170240 | 00:00:58 |
| 40.0.3.138 | 0 | 2660 | 170240 | 00:00:58 |

| Packet count for last active timeout | Byte count for last active timeout |
|--------------------------------------|------------------------------------|
| 2805 | 179520 |
| 2805 | 179520 |

show services accounting flow-detail limit

In this sample, the output is split into three sections, with ellipses (...) indicating where the sections are continued.

```
user@host> show services accounting flow-detail limit 1
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
Protocol   Input          Source          Source   Output
           interface    address         port     interface...
tcp(6)     ge-5/0/1.0    11.1.1.2        0        ge-5/0/0.0

Destination      Destination      Packet   Byte   Time since last
address          port            count    count active timeout...
40.0.3.149              0            2158    138112      00:00:47

Packet count for   Byte count for
last active timeout last active timeout
                2827                180928
```

show services accounting flow-detail name extensive

```
user@host> show services accounting flow-detail name cf-2 extensive
Service Accounting interface: mo-0/2/0, Local interface index: 145
Service name: cf-2
  TOS: 0, Protocol: udp(17), TCP flags: 0
  Source address: 10.10.10.1, Source prefix length: 0, Destination address:
20.20.20.20,
  Destination prefix length: 0, Source port: 1173, Destination port: 69
  Input SNMP interface index: 65, Output SNMP interface index: 0, Source-AS: 0,
  Destination-AS: 0
  Start time: 62425, End time: 635265, Packet count: 165845, Byte count: 9453165
```

show services accounting flow-detail limit order bytes

The output of the following command is displayed over 141 columns, not the standard 80 columns. In this sample, the output is split into three sections, with ellipses (...) indicating where the sections are continued.

```
user@host> show services accounting flow-detail limit 5 order bytes
Service Accounting interface: mo-2/0/0, Local interface index: 356
Service name: (default sampling)
Protocol   Input          Source          Source   Output
           interface    address         port     interface...
icmp(1)    ge-2/3/0.0    11.1.1.2        0        .local.
icmp(1)    ge-2/3/0.0    11.1.1.2        0        .local.
icmp(1)    ge-2/3/0.0    11.1.1.2        0        .local.
icmp(1)    ge-2/3/0.0    11.1.1.2        0        .local.
icmp(1)    ge-2/3/0.0    11.1.1.2        0        .local.

Destination      Destination      Packet   Byte   Time since last
address          port            count    count active timeout...
51.88.128.2              0            16      12148      Not applicable
52.78.144.2              0            16      15229      Not applicable
51.147.192.2             0            16      13296      Not applicable
51.136.16.2              0            16      13924      Not applicable
50.214.48.2              0            16      13428      Not applicable

Packet count for   Byte count for
```

| | |
|---------------------|---------------------|
| last active timeout | last active timeout |
| Not applicable | Not applicable |
| Not applicable | Not applicable |
| Not applicable | Not applicable |
| Not applicable | Not applicable |
| Not applicable | Not applicable |

show services accounting flow-detail source-port

```
user@host> show services accounting flow-detail name cf-2 detail source-port 1173
Service Accounting interface: mo-0/2/0, Local interface index: 145
Service name: cf-2
  Protocol: udp(17), Source address: 10.10.10.1, Source port: 1173, Destination
address:
20.20.20.20, Destination port: 69
  Start time: 62425, End time: 811115, Packet count: 142438, Byte count: 8118966
```


show services accounting memory

| | |
|---------------------------------|---|
| Syntax | show services accounting memory |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display memory and flow record statistics. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show services accounting memory (Monitoring PIC interface) on page 2223
show services accounting memory (Service PIC interface) on page 2224 |
| Output Fields | Table 136 on page 2223 lists the output fields for the show services accounting memory command. Output fields are listed in the approximate order in which they appear. |

Table 136: show services accounting memory Output Fields

| Output Field | Output Field Description |
|------------------------------|---|
| Service Accounting interface | Name of the service accounting interface. |
| Memory Utilization | |
| Local interface index | Index counter of the local interface. |
| Allocation count | Number of flow records allocated. |
| Free count | Number of flow records freed. |
| Maximum allocated | Maximum number of flow records allocated since the monitoring station booted. This number represents the peak number of flow records allocated at a time. |
| Allocations per second | Flow records allocated per second during the last statistics interval on the PIC. |
| Frees per second | Flow records freed per second during the last statistics interval on the PIC. |
| Total memory used | Total amount of memory currently used (in bytes). |
| Total memory free | Total amount of memory currently free (in bytes). |

Sample Output

show services accounting memory (Monitoring PIC interface)

```

user@host> show services accounting memory
Service Accounting interface: mo-2/0/0, Local interface index: 468
Memory utilization

```

```
Allocation count: 437340, Free count: 433699, Maximum allocated: 6782
Allocations per second: 3366, Frees per second: 6412
Total memory used (in bytes): 133460320,
Total memory free (in bytes): 133918352
```

Sample Output

show services accounting memory (Service PIC interface)

```
user@host> show services accounting memory
Service Accounting interface: sp-0/1/0
Memory utilization
  Allocation count: 1000, Free count: 0
  Allocations per second: 0, Frees per second: 0
  Total memory used (in bytes): 218158272
  Total memory free (in bytes): 587147696

Service Accounting interface: sp-1/0/0
Memory utilization
  Allocation count: 1000, Free count: 0
  Allocations per second: 0, Frees per second: 0
  Total memory used (in bytes): 218157592
  Total memory free (in bytes): 587148376
```

show services accounting packet-size-distribution

| | |
|---------------------------------|--|
| Syntax | show services accounting packet-size-distribution
<name (* all <i>service-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display a packet size distribution histogram. |
| Options | <p>none—Display a packet size distribution histogram of all accounting services.</p> <p>name (* all <i>service-name</i>)—(Optional) Display a packet size distribution histogram. Use a wildcard character, specify all services, or provide a specific services name.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services accounting packet-size-distribution name on page 2225 |
| Output Fields | Table 137 on page 2225 lists the output fields for the show services accounting packet-size-distribution command. Output fields are listed in the approximate order in which they appear. |

Table 137: show services accounting packet-size-distribution Output Fields

| Field Name | Field Description |
|------------------------------|---|
| Service Accounting interface | Name of the service accounting interface. |
| Service name | Name of a service that was configured at the [edit-forwarding-options accounting] hierarchy level. The default display, (default sampling), indicates the service was configured at the [edit-forwarding-options sampling-level] hierarchy level. |
| Local interface index | Index counter of the local interface. |
| Range start | Smallest packet length (in bytes) to count. |
| Range end | Largest packet length (in bytes) to count. |
| Number of packets | Count of packets detected in the size between Range start and Range end. |
| Percentage packets | Percentage of the total number of packets that are in this size range. |

Sample Output

show services accounting packet-size-distribution name

```
user@host> show services accounting packet-size-distribution name test3
Service Accounting interface: mo-0/2/0, Local interface index: 163
Service name: test3
```

| Range start | Range end | Number of packets | Percentage packets |
|-------------|-----------|-------------------|--------------------|
| 32 | 64 | 2924 | 100 |

show services accounting status

| | |
|---------------------------------|--|
| Syntax | <code>show services accounting status</code>
<code><inline-jflow fpc-slot <i>slot-number</i> name (* all <i>service-name</i>)></code> |
| Release Information | Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 13.2R2 for EX Series switches. |
| Description | Display available Physical Interface Cards (PICs) for accounting services. |
| Options | <p>none—Display available PICs for all accounting services.</p> <p>inline-jflow fpc-slot <i>slot-number</i>—(Optional) Display inline flow accounting status for the specified FPC. For a two-member MX Series Virtual Chassis or EX9200 Virtual Chassis, the master router or switch uses FPC slot numbers 0 through 11 with no offset; the backup router or switch uses FPC slot numbers 12 through 23, with an offset of 12.</p> <p>name (* all <i>service-name</i>)—(Optional) Display available PICs. Use a wildcard character, specify all services, or provide a specific services name.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> show services accounting flow on page 2213 Inline Flow Monitoring for Virtual Chassis Overview |
| List of Sample Output | <p>show services accounting status name (Monitoring PIC interface) on page 2228</p> <p>show services accounting status name (Service PIC interface) on page 2228</p> <p>show services accounting status inline-jflow fpc-slot <i>slot-number</i> (when both IPv4 and IPv6 are configured) on page 2229</p> <p>show services accounting status inline-jflow (MX80 Router when both IPv4 and IPv6 are configured) on page 2229</p> |
| Output Fields | Table 138 on page 2227 lists the output fields for the show services accounting status command. Output fields are listed in the approximate order in which they appear. |

Table 138: show services accounting status Output Fields

| Field | Field Description |
|------------------------------|---|
| Service Accounting interface | Name of the service accounting interface. |
| Service name | Name of a service that was configured at the <code>[edit-forwarding-options accounting]</code> hierarchy level. The default display, <code>(default sampling)</code> , indicates the service was configured at the <code>[edit-forwarding-options sampling-level]</code> hierarchy level. |
| FPC Slot | Slot number of the FPC for which the flow information is displayed. |

Table 138: show services accounting status Output Fields (*continued*)

| Field | Field Description |
|------------------------------|---|
| Local interface index | Index counter of the local interface. |
| Interface state | Accounting state of the passive monitoring interface. <ul style="list-style-type: none"> • Accounting—PIC is actively accounting. • Disabled—PIC has been disabled from the CLI. • Not accounting—PIC is up but not accounting. This can happen while the PIC is coming online, or when the PIC is up but has no logical unit configured under the physical interface. • Unknown |
| Group index | Integer that represents the monitoring group of which the PIC is a member. Group index is a mapping from the group name to an index. It is not related to the number of monitoring groups. |
| Export interval (in seconds) | Configured export interval for cflowd records, in seconds. |
| Export format | Configured export format. |
| Protocol | Protocol the PIC is configured to monitor. |
| Engine type | Configured engine type that is inserted in output cflowd packets. |
| Engine ID | Configured engine ID that is inserted in output cflowd packets. |
| Route Record Count | Number of routes recorded. |
| AS Record Count | Number of autonomous systems recorded. |
| Route Records Set | Status of route recording; whether routes are recorded or not. |
| Configuration Set | Status of monitoring configuration; whether monitoring configuration is set or not. |

Sample Output

show services accounting status name (Monitoring PIC interface)

```

user@host> show services accounting status name count1
Service Accounting interface: mo-2/0/0, Local interface index: 468
Service name: count1
Interface state: Accounting
  Group index: 0
  Export interval (in seconds): 60, Export format: cflowd v8
  Protocol: IPv4, Engine type: 55, Engine ID: 5

```

Sample Output

show services accounting status name (Service PIC interface)

```

user@host> show services accounting status name

```

```

Service Accounting interface: sp-0/1/0
Interface state: Accounting
  Export format: 9, Route record count: 0
  IFL to SNMP index count: 7, AS count: 0
  Configuration set: Yes, Route record set: No, IFL SNMP map set: Yes

Service Accounting interface: sp-1/0/0
Interface state: Accounting
  Export format: 9, Route record count: 33
  IFL to SNMP index count: 7, AS count: 1
  Configuration set: Yes, Route record set: Yes, IFL SNMP map set: Yes

```

show services accounting status inline-jflow fpc-slot slot-number (when both IPv4 and IPv6 are configured)

```

user@host> show services accounting status inline-jflow fpc-slot 5
FPC Slot: 5
  IPv4 export format: Version-IPFIX, IPv6 export format: Version-IPFIX
  VPLS export format: Not set
  IPv4 Route Record Count: 5, IPv6 Route Record Count: 7
  Route Record Count: 12, AS Record Count: 1
  Route-Records Set: Yes, Config Set: Yes

```

show services accounting status inline-jflow (MX80 Router when both IPv4 and IPv6 are configured)

```

user@host> show services accounting status inline-jflow

Status information
  TFEB Slot: 0
  Export format: IP-FIX
  IPv4 Route Record Count: 6, IPv6 Route Record Count: 8
  Route Record Count: 14, AS Record Count: 1
  Route-Records Set: Yes, Config Set: Yes

```

show services accounting usage

| | |
|---------------------------------|---|
| Syntax | show services accounting usage
<name <i>service-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display the CPU usage of PIC used for active flow monitoring. |
| Options | <p>none—Display CPU usage for all service names.</p> <p>name <i>service-name</i>—(Optional) Display CPU usage for the specified service name.</p> |
| Additional Information | When no route record has been downloaded from the PIC, this command reports no flows, even though packets are being sampled. |
| Required Privilege Level | view |
| List of Sample Output | show services accounting usage (Monitoring PIC interface) on page 2231
show services accounting usage (Service PIC interface) on page 2231 |
| Output Fields | Table 139 on page 2230 lists the output fields for the show services accounting usage command. Output fields are listed in the approximate order in which they appear. |

Table 139: show services accounting usage Output Fields

| Output Field | Output Field Description |
|------------------------------|--|
| Service Accounting interface | Name of the service accounting interface. |
| Service name | Name of a service that was configured at the [edit-forwarding-options accounting] hierarchy level. The default display, (default sampling), indicates the service was configured at the [edit-forwarding-options sampling-level] hierarchy level. |
| Local interface index | Index counter of the local interface. |
| Uptime | Time that the PIC has been operational (in milliseconds). |
| Interrupt time | Total time that the PIC has spent processing packets since the last PIC reset (in microseconds). |
| Load (5 second) | CPU load on the PIC, averaged more than 5 seconds. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time. |
| Load (1 minute) | CPU load on the PIC, averaged more than 1 minute. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time. |

Sample Output

show services accounting usage (Monitoring PIC interface)

```
user@host> show services accounting usage
Service Accounting interface: mo-1/1/0, Local interface index: 15
Service name: (default sampling)
CPU utilization
  Uptime: 600413856 milliseconds, Interrupt time: 2403 microseconds
  Load (5 second): 43%, Load (1 minute): 24%
```

Sample Output

show services accounting usage (Service PIC interface)

```
user@host> show services accounting usage
Service Accounting interface: sp-0/1/0
Service name: (default sampling)
CPU utilization
  Uptime: 7853940 milliseconds, Interrupt time: 0 microseconds
  Load (5 second): 2%, Load (1 minute): 0%

Service Accounting interface: sp-0/1/0
Service name: (default sampling)
CPU utilization
  Uptime: 331160 milliseconds, Interrupt time: 0 microseconds
  Load (5 second): 2%, Load (1 minute): 0%
```

show services dynamic-flow-capture content-destination

| | |
|---------------------------------|--|
| Syntax | show services dynamic-flow-capture content-destination capture-group <i>group-name</i>
destination-identifier <i>identifier</i>
<terse> |
| Release Information | Command introduced in Junos OS Release 7.4. |
| Description | (M320 routers and T Series routers only) Display information about the content destination that receives packets from the dynamic flow capture (DFC) interface. |
| Options | <p>capture-group <i>group-name</i>—Capture-group identifier.</p> <p>destination-identifier <i>identifier</i>—Content destination identifier.</p> <p>terse—(Optional) Display summary information.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services dynamic-flow-capture content-destination on page 2233 |
| Output Fields | Table 140 on page 2232 lists the output fields for the show services dynamic-flow-capture content-destination command. Output fields are listed in the approximate order in which they appear. |

Table 140: show services dynamic-flow-capture content-destination Output Fields

| Output Field | Output Field Description | Level of Output |
|---------------------------------|--|-----------------|
| Capture group | Name of the capture group. | to be provided |
| Content destination | Name of the content destination. | to be provided |
| Criteria | Number of criteria specified. | to be provided |
| Bandwidth | Bandwidth used by the matched traffic. | to be provided |
| Matched packets | Number of matched packets sent to the content destination. | to be provided |
| Matched bytes | Number of matched bytes sent to the content destination. | to be provided |
| Congestion notifications | Number of notification messages sent. | to be provided |

Sample Output

`show services dynamic-flow-capture content-destination`

```
user@host> show services dynamic-flow-capture content-destination capture-group g1
destination-identifier cd1 terse
  Capture group: g1, Content destination: cd1, Criteria: 0, Bandwidth: 0, Matched
  packets: 0, Matched bytes: 0, Congestion notifications: 0
```

show services dynamic-flow-capture control-source

| | |
|---------------------------------|--|
| Syntax | <code>show services dynamic-flow-capture control-source capture-group <i>group-name</i>
control-source <i>identifier</i>
<detail terse></code> |
| Release Information | Command introduced in Junos OS Release 7.4. |
| Description | (M320 routers and T Series routers only) Display information about the control source that makes dynamic flow capture requests to the dynamic flow capture interface. |
| Options | <p><code>capture-group <i>group-name</i></code>—Capture group identifier.</p> <p><code>control-source <i>identifier</i></code>—Control source identifier.</p> <p><code>detail terse</code>—(Optional) Display the specified level of output.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services dynamic-flow-capture control-source on page 2235
show services dynamic-flow-capture control-source detail on page 2235 |
| Output Fields | Table 141 on page 2234 lists the output fields for the <code>show services dynamic-flow-capture control-source</code> command. Output fields are listed in the approximate order in which they appear. |

Table 141: show services dynamic-flow-capture control-source Output Fields

| Output Field | Output Field Description |
|-------------------------------------|--|
| Capture group | Name of the capture group. |
| Control source | Name of the control source. |
| Criteria added, Criteria add failed | Number of criteria added or added and failed. |
| Active criteria | Number of active criteria. |
| Static criteria, Dynamic criteria | Number of static or dynamic criteria. |
| Control protocol requests | Total number of control protocol requests. |
| Requests | Number of Add , Delete , List , Refresh , and No-op control protocol requests. |
| Failed | Number of Add , Delete , List , Refresh , and No-op failed control protocol requests. |
| Add request rate | Rate of add requests. |

Table 141: show services dynamic-flow-capture control-source Output Fields (*continued*)

| Output Field | Output Field Description |
|-------------------------------|--|
| Add request peak rate | Peak rate of add requests. |
| Bandwidth across all criteria | Bandwidth used by all the requests. |
| Total notifications | Total number of notifications sent and the number of notifications by category: Restart , Rollover , Timeout , Congestion , Congestion delete , and Dups (duplicates) dropped. |
| Criteria deleted | Total number of criteria deleted and the number of deleted criteria by category: Timeout idle , Timeout total , Packets , and Bytes . |
| Sequence number | Sequence number. |

Sample Output

show services dynamic-flow-capture control-source

```

user@host> show services dynamic-flow-capture control-source source-identifier cs0_cg0
capture-group cg_0
Capture group: cg_0, Control source: cs0_cg0
Criteria added: 28, Criteria add failed: 0, Active criteria: 0, Control protocol
requests: 28, Add request rate: 0,
Add request peak rate: 1, Bandwidth across all criteria: 0, Total notifications:
1, Criteria deleted: 28, Sequence number: 0

```

show services dynamic-flow-capture control-source detail

```

user@host> show services dynamic-flow-capture control-source source-identifier cs0_cg0
capture-group cg_0 detail
Capture group: cg_0, Control source: cs0_cg0
Criteria added: 28, Criteria add failed: 0
Active criteria: 0
Static criteria: 0, Dynamic criteria: 0
Control protocol requests: 28

```

| | Add | Delete | List | Refresh | No-op |
|----------|-----|--------|------|---------|-------|
| Requests | 28 | 0 | 0 | 0 | 0 |
| Failed | 0 | 0 | 0 | 0 | 0 |

```

Add request rate: 0
Add request peak rate: 1
Bandwidth across all criteria: 0
Total notifications: 1
Restart: 1, Rollover: 0, No-op: 0, Timeout: 0, Congestion: 0, Congestion
delete: 0, Dups dropped: 0
Criteria deleted: 28
Timeout idle: 0, Timeout total: 0, Packets: 0, Bytes: 0
Sequence number: 0

```

show services dynamic-flow-capture statistics

| | |
|---------------------------------|--|
| Syntax | <code>show services dynamic-flow-capture statistics capture-group <i>group-name</i></code> |
| Release Information | Command introduced in Junos OS Release 7.4. |
| Description | (M320 routers and T Series routers only) Display statistics information about the capture group specified for dynamic flow capture. |
| Options | <code>capture-group <i>group-name</i></code> —Capture group identifier. |
| Required Privilege Level | view |
| List of Sample Output | show services dynamic-flow-capture statistics on page 2237 |
| Output Fields | Table 142 on page 2236 lists the output fields for the show services dynamic-flow-capture statistics command. Output fields are listed in the approximate order in which they appear. |

Table 142: show services dynamic-flow-capture statistics Output Fields

| Output Field | Output Field Description |
|------------------------|--|
| Input | <p>Incoming dynamic flow capture packet statistics:</p> <ul style="list-style-type: none"> • Control protocol packets—Number of control protocol packets received. • Captured data packets—Number of data packets captured. • Control IRI packets—Number of control IRI packets received. |
| Control protocol drops | <p>Control protocol packets dropped for the following reasons:</p> <ul style="list-style-type: none"> • Not IP packets—Dropped packets were not IP packets. • Not UDP packets—Dropped packets were not User Datagram Protocol (UDP) packets. • Invalid destination address—Dropped packets had invalid destination addresses. • No memory—Packets dropped because of insufficient memory. • Unauthorized control source—Packets dropped because the control source was not authenticated. • Bad request—Packets dropped because the request was invalid. • Unknown control source—Packets dropped because the control source was not known. • Not DTCP—Dropped packets did not adhere to the control protocol format. • Bad command line—Packets dropped because of a version mismatch. • Bandwidth exceeded—Packets dropped because the bandwidth was exceeded. • Drop rate due to exceeded bandwidth—Rate of traffic dropped because the bandwidth was exceeded. • Other—Packets dropped for other reasons or undetermined causes. |

Table 142: show services dynamic-flow-capture statistics Output Fields (*continued*)

| Output Field | Output Field Description |
|------------------------|--|
| Input drops | <p>Incoming dynamic flow capture packets dropped for the following reasons:</p> <ul style="list-style-type: none"> • Unknown packets—Packets dropped because the packet type was not recognized. • Captured data not IPv4—Packets dropped because they were not IPv4 packets. • Captured data too small—Packets dropped because they were smaller than the size reported in their headers. • Captured data drops—Data packets dropped because of undetermined causes. • Captured data not matched—Packets dropped because they did not match filter criteria. • Bandwidth exceeded—Packets dropped because the bandwidth was exceeded. • Drop rate due to exceeded bandwidth—Rate of traffic dropped because the bandwidth was exceeded. |
| Output | <p>Outgoing dynamic flow capture packet statistics:</p> <ul style="list-style-type: none"> • Control protocol packets—Number of control protocol packets sent. • Captured data packets—Number of captured data packets sent. |
| Output drops | <p>Outgoing packets dropped:</p> <ul style="list-style-type: none"> • Control protocol drops—Number of control protocol packets dropped. • Captured data drops—Number of captured data packets dropped. |
| Flow Statistics | <p>DFC flow statistics:</p> <ul style="list-style-type: none"> • Active flow cache entries • Active flow cache usage percentage • Flow cache entries allocated • Number of control sources • Number of content destinations • Number of criteria • Maximum criteria matching one flow • Cached flows purged for memory • Maximum filters matching one packet |

Sample Output

show services dynamic-flow-capture statistics

```

user@host> show services dynamic-flow-capture statistics capture-group g1
Input:

  Control protocol packets: 643, Captured data packets: 69977, Control IRI packets:
  337

Control protocol drops:

  Not IP packets: 0, Not UDP packets: 3, Invalid destination address: 0, No memory:
  0, Unauthorized control source: 0,

  Bad request: 0, Unknown control source: 0, Not DTCP: 0, Bad command line: 0,
  Bandwidth exceeded: 0,

```

Drop rate due to exceeded bandwidth: 0, Other: 0

Input drops:

Unknown packets: 0, Captured data not IPv4: 0, Captured data too small: 0,
Captured data drops: 0, Captured data not matched: 0,

Bandwidth exceeded: 0, Drop rate due to exceeded bandwidth: 0

Output:

Control protocol packets: 644, Captured data packets: 1119624

Output drops:

Control protocol drops: 0, Captured data drops: 0

Flow Statistics:

Active flow cache entries: 40, Active flow cache usage percentage: 0, Flow cache
entries allocated: 40,

Number of control sources: 4, Number of content destinations: 64, Number of
criteria: 640,

Maximum criteria matching one flow: 16, Cached flows purged for memory: 0,
Maximum filters matching one packet: 16

show services flow-collector file interface

| | |
|---------------------------------|--|
| Syntax | show services flow-collector file interface (all cp-fpc/pic/port)
<detail extensive terse> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M40e, M160, and M320 routers and T Series routers only) Display information about flow collector files. |
| Options | <p>all cp-fpc/pic/port—Display file information for all configured flow collector interfaces or for the specified interface.</p> <p>detail extensive terse—(Optional) Display the specified level of output.</p> |
| Additional Information | No entries are displayed for files that have been successfully transferred. |
| Required Privilege Level | view |
| List of Sample Output | show services flow-collector file interface extensive on page 2240 |
| Output Fields | Table 143 on page 2239 lists the output fields for the show services flow-collector file interface command. Output fields are listed in the approximate order in which they appear. |

Table 143: show services flow-collector file interface Output Fields

| Output Field | Output Field Description | Level of Output |
|-------------------|---|------------------|
| Filename | Name of the file created on the flow collector interface. | All levels |
| Flows | Total number of collector flows for which records are present in the file. | none specified |
| Throughput | Throughput statistics: <ul style="list-style-type: none"> • Flow records—Number of flow records in the file. <ul style="list-style-type: none"> • per second—Average number of flow records per second. • peak per second—Peak number of flow records per second. • Uncompressed bytes—Total file size before compression. <ul style="list-style-type: none"> • per second—Average number of uncompressed bytes per second. • peak per second—Peak number of uncompressed bytes per second. • Compressed bytes—Total file size after compression. <ul style="list-style-type: none"> • per second—Average number of compressed bytes per second. • peak per second—Peak number of compressed bytes per second. | extensive |

Table 143: show services flow-collector file interface Output Fields (*continued*)

| Output Field | Output Field Description | Level of Output |
|--------------|---|-----------------|
| Status | <p>File statistics:</p> <ul style="list-style-type: none"> • Compressed blocks—(extensive output only) Data blocks in the file that have been compressed. The file is exported only when the compressed block count and block count become the same. • Block count—(extensive output only) Total number of data blocks in the file. • State—Processing state of the file. <ul style="list-style-type: none"> • Active—The flow collector interface is writing to the file. • Export 1—File export is in progress to the primary server. • Export 2—File export is in progress to the secondary server. • Wait—File is pending export. • Transfer attempts 0—Number of attempts made to transfer the file. If the file is successfully transferred in the first attempt, this field is 0. | All levels |

Sample Output

show services flow-collector file interface extensive

```

user@host> show services flow-collector file interface cp-3/2/0 extensive
Filename: cFlowd-py69Ni69-0-20031112_014301-so_3_0_0_0.bcp.bi.gz
Throughput:
  Flow records: 188365, per second: 238, peak per second: 287
  Uncompressed bytes: 21267756, per second: 27007, peak per second: 32526
  Compressed bytes: 2965643, per second: 0, peak per second: 22999
Status:
  Compressed blocks: 156, Block count: 156
  State: Active, Transfer attempts: 0

```

show services flow-collector input interface

| | |
|---------------------------------|--|
| Syntax | show services flow-collector input interface (all cp-fpc/pic/port)
<detail extensive terse> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M40e, M160, and M320 routers and T Series routers only) Display the number of packets received by collector interfaces from monitoring interfaces. |
| Options | <p>all cp-fpc/pic/port—Display packets received by all configured flow collector interfaces or by the specified interface.</p> <p>detail extensive terse—(Optional) Display the specified level of output.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services flow-collector input interface on page 2241
show services flow-collector input interface all on page 2241 |
| Output Fields | Table 144 on page 2241 lists the output fields for the show services flow-collector input interface command. Output fields are listed in the approximate order in which they appear. |

Table 144: show services flow-collector input interface Output Fields

| Output Field | Output Field Description |
|------------------|--|
| Interface | Name of the monitoring interface. |
| Packets | Number of packets traveling from the monitoring interface to the flow collector interface. |
| Bytes | Number of bytes traveling from the monitoring interface to the flow collector interface. |

Sample Output

show services flow-collector input interface

```

user@host> show services flow-collector input interface cp-3/2/0
Interface                Packets    Bytes
mo-3/0/0.0               21706     32328568
mo-3/1/0.0               21706     32329096

```

show services flow-collector input interface all

```

user@host> show services flow-collector input interface all
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
Interface                Packets    Bytes
mo-3/0/0.0               274        416232
mo-3/3/0.0               274        416184

```

| | | |
|------------|-----|--------|
| mo-1/0/0.0 | 274 | 416232 |
| mo-1/1/0.0 | 274 | 416232 |
| mo-1/2/0.0 | 274 | 416232 |
| mo-1/3/0.0 | 274 | 416232 |
| mo-3/1/0.0 | 274 | 416232 |
| mo-4/0/0.0 | 274 | 416232 |
| mo-4/1/0.0 | 274 | 416232 |
| mo-4/2/0.0 | 274 | 416184 |
| mo-4/3/0.0 | 274 | 416232 |
| mo-5/0/0.0 | 274 | 416232 |
| mo-5/1/0.0 | 274 | 416232 |
| mo-5/2/0.0 | 274 | 416232 |
| mo-5/3/0.0 | 274 | 416232 |
| mo-6/0/0.0 | 274 | 416232 |

Flow collector interface: cp-6/3/0
Interface state: Collecting flows

show services flow-collector interface

| | |
|---------------------------------|---|
| Syntax | show services flow-collector interface (all cp-fpc/pic/port)
<detail extensive terse> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M40e, M160, and M320 routers and T Series routers only) Display overall statistics for the flow collector application. |
| Options | <p>all cp-fpc/pic/port—Display statistics for flow collector applications on all interfaces or for the specified interface.</p> <p>detail extensive terse—(Optional) Display the specified level of output.</p> |
| Required Privilege Level | view |
| List of Sample Output | show services flow-collector interface all detail on page 2245
show services flow-collector interface all extensive on page 2246
show services flow-collector interface all terse on page 2248
show services flow-collector interface extensive on page 2248 |
| Output Fields | Table 145 on page 2243 lists the output fields for the show services flow-collector interface command. Output fields are listed in the approximate order in which they appear. |

Table 145: show services flow-collector interface Output Fields

| Output Field | Output Field Description | Level of Output |
|--------------------------|---|------------------|
| Flow collector interface | Name of the flow collector interface. | All levels |
| Interface state | Collecting flow state for the interface. | All levels |
| Packets | Total number of packets received. | none specified |
| Flows Uncompressed Bytes | Total uncompressed data size for all files created on this PIC. | none specified |
| Compressed Bytes | Total compressed data size for all files created on this PIC. | none specified |
| FTP bytes | Total number of bytes transferred to the FTP server, including those dropped during transfer. | none specified |
| FTP files | Total number of FTP transfers attempted by the server. | none specified |
| Memory | Bytes used on the PIC and bytes free. | detail extensive |

Table 145: show services flow-collector interface Output Fields (*continued*)

| Output Field | Output Field Description | Level of Output |
|-------------------|---|-------------------------|
| Input | Incoming flow collector packet statistics: <ul style="list-style-type: none"> • Packets—Number of packets received on the unit. <ul style="list-style-type: none"> • per second—Average number of packets per second. • peak per second—Peak number of packets per second. • Bytes—Number of bytes received on the unit. <ul style="list-style-type: none"> • per second—Average number of bytes per second. • peak per second—Peak number of bytes per second. • Flow records processed—Number of records in the flow collector packets that were processed by the flow-collector interface. <ul style="list-style-type: none"> • per second—Average number of flow records processed per second. • peak per second—Peak number of flow records per second. | detail extensive |
| Allocation | Data block statistics: <ul style="list-style-type: none"> • Blocks allocated—Total number of data blocks (containing flow records) allocated to the files created on this PIC. <ul style="list-style-type: none"> • per second—Average number of blocks allocated per second. • peak per second—Peak number of blocks allocated per second. • Blocks freed—Total number of data blocks freed. <ul style="list-style-type: none"> • per second—Average number of blocks freed per second. • peak per second—Peak number of blocks freed per second. • Blocks unavailable—Total number of data block requests denied, typically because of a memory shortage. <ul style="list-style-type: none"> • per second—Average number of blocks unavailable per second. • peak per second—Peak number of blocks unavailable per second. | extensive |
| Files | File statistics, incremented since the PIC last booted: <ul style="list-style-type: none"> • Files created—Total number of files created on this PIC. • Files exported— Number of files successfully created and exported. • Files destroyed— (extensive output only) Number of files successfully exported and files dropped by the flow collection interface. | detail extensive |
| Throughput | Throughput statistics: <ul style="list-style-type: none"> • Uncompressed bytes—Total uncompressed data size for all files created on this PIC. <ul style="list-style-type: none"> • per second—Average number of uncompressed bytes per second. • peak per second—Peak number of uncompressed bytes per second. • Compressed bytes—Total compressed data size for all files created on this PIC. <ul style="list-style-type: none"> • per second—Average number of compressed bytes per second. • peak per second—Peak number of compressed bytes per second. | detail extensive |

Table 145: show services flow-collector interface Output Fields (*continued*)

| Output Field | Output Field Description | Level of Output |
|---------------------------------|--|-------------------------|
| Packet drops | <p>Number of packets dropped for the following causes:</p> <ul style="list-style-type: none"> • No memory—Packets dropped because of insufficient memory. • Not IP—Packets dropped because they are not IP packets. • Not IPv4—Packets dropped because they are not IP version 4 packets. • Too small—Packets dropped because each packet was smaller than the size reported in its header. • Fragments—Packets dropped because of fragmentation. Fragments are not reassembled. • ICMP—Packets dropped because they are not ICMP packets. • TCP—Packets dropped because they are not TCP packets. • Unknown—Packets dropped because of undetermined causes. • Not Junos flow—Packets dropped because they are not interpreted by Junos OS. Junos OS interprets only IPv4, UDP cflowd version 5 packets. | extensive |
| File transfer | <p>File transfer statistics:</p> <ul style="list-style-type: none"> • FTP bytes—Total number of bytes transferred to the FTP server, including those dropped during transfer. • FTP files—Total number of FTP transfers attempted by the server. • FTP failure—Total number of FTP failures encountered by the server. | detail extensive |
| Flow collector interface | Physical interface acting as a flow collector. | detail |
| Export channel | <p>Export channel 0 is unit 0. Export channel 1 is unit 1. Flow receive channel is unit 2. Server status statistics are the following:</p> <ul style="list-style-type: none"> • Current server Primary or Secondary—Current FTP server being used. Value is • Primary server state—State of the server: <ul style="list-style-type: none"> • OK—Server is operating without problems. • FTP error—Server encountered an FTP protocol error while sending files. • Network error—Flow-collector interface has errors when contacting the primary FTP server. • Unknown—First file transfer has not been sent to the primary server. • Secondary server state—State of the server: <ul style="list-style-type: none"> • OK—Server is operating without errors. • FTP error—Server encountered an FTP protocol error while sending files. • Network error—Flow-collector interface has errors when contacting the secondary FTP server. • Unknown—First file transfer has not been sent to the secondary server. • Not configured—Secondary server is not configured. | detail extensive |

Sample Output

show services flow-collector interface all detail

```
user@host> show services flow-collector interface all detail
```

```
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
Memory:
  Used: 51452732, Free: 440329088
Input:
  Packets: 4384, per second: 0, peak per second: 156
  Bytes: 6659616, per second: 0, peak per second: 249695
  Flow records processed: 131070, per second: 0, peak per second: 4914
Files:
  Files created: 1, per second: 0, peak per second: 0
  Files exported: 1, per second: 0, peak per second: 0
Throughput:
  Uncompressed bytes: 13742307, per second: 0, peak per second: 593564
  Compressed bytes: 3786177, per second: 0, peak per second: 162826
File Transfer:
  FTP bytes: 3786247, per second: 0, peak per second: 378620
  FTP files: 1, per second: 0, peak per second: 0
  FTP failure: 0
Export channel: 0
  Current server: Primary
  Primary server state: OK, Secondary server state: OK
Export channel: 1
  Current server: Primary
  Primary server state: Unknown, Secondary server state: OK
```

```
Flow collector interface: cp-6/3/0
Interface state: Collecting flows
Memory:
  Used: 51452732, Free: 440329088
Input:
  Packets: 0, per second: 0, peak per second: 0
  Bytes: 0, per second: 0, peak per second: 0
  Flow records processed: 0, per second: 0, peak per second: 0
Files:
  Files created: 0, per second: 0, peak per second: 0
  Files exported: 0, per second: 0, peak per second: 0
Throughput:
  Uncompressed bytes: 0, per second: 0, peak per second: 0
  Compressed bytes: 0, per second: 0, peak per second: 0
File Transfer:
  FTP bytes: 70, per second: 0, peak per second: 6
  FTP files: 0, per second: 0, peak per second: 0
  FTP failure: 0
Export channel: 0
  Current server: Primary
  Primary server state: Unknown, Secondary server state: OK
Export channel: 1
  Current server: Primary
  Primary server state: Unknown, Secondary server state: OK
```

show services flow-collector interface all extensive

```
user@host> show services flow-collector interface all extensive
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
Memory:
  Used: 51452732, Free: 440329088
Input:
  Packets: 4384, per second: 0, peak per second: 156
  Bytes: 6659616, per second: 0, peak per second: 249695
  Flow records processed: 131070, per second: 0, peak per second: 4914
```



```

Allocation:
  Blocks allocated: 108, per second: 0, peak per second: 0
  Blocks freed: 108, per second: 0, peak per second: 10
  Blocks unavailable: 0, per second: 0, peak per second: 0
Files:
  Files created: 1, per second: 0, peak per second: 0
  Files exported: 1, per second: 0, peak per second: 0
  Files destroyed: 1, per second: 0, peak per second: 0
Throughput:
  Uncompressed bytes: 13742307, per second: 0, peak per second: 593564
  Compressed bytes: 3786177, per second: 0, peak per second: 162826
Packet drops:
  No memory: 0, Not IP: 0
  Not IPv4: 0, Too small: 0
  Fragments: 0, ICMP: 0
  TCP: 0, Unknown: 0
  Not JUNOS flow: 0
File Transfer:
  FTP bytes: 3786247, per second: 0, peak per second: 378620
  FTP files: 1, per second: 0, peak per second: 0
  FTP failure: 0
Export channel: 0
  Current server: Primary
  Primary server state: OK, Secondary server state: OK
Export channel: 1
  Current server: Primary
  Primary server state: Unknown, Secondary server state: OK

Flow collector interface: cp-6/3/0
Interface state: Collecting flows
Memory:
  Used: 51452732, Free: 440329088
Input:
  Packets: 0, per second: 0, peak per second: 0
  Bytes: 0, per second: 0, peak per second: 0
  Flow records processed: 0, per second: 0, peak per second: 0
Allocation:
  Blocks allocated: 0, per second: 0, peak per second: 0
  Blocks freed: 0, per second: 0, peak per second: 0
  Blocks unavailable: 0, per second: 0, peak per second: 0
Files:
  Files created: 0, per second: 0, peak per second: 0
  Files exported: 0, per second: 0, peak per second: 0
  Files destroyed: 0, per second: 0, peak per second: 0
Throughput:
  Uncompressed bytes: 0, per second: 0, peak per second: 0
  Compressed bytes: 0, per second: 0, peak per second: 0
Packet drops:
  No memory: 0, Not IP: 0
  Not IPv4: 0, Too small: 0
  Fragments: 0, ICMP: 0
  TCP: 0, Unknown: 0
  Not JUNOS flow: 0
File Transfer:
  FTP bytes: 70, per second: 0, peak per second: 6
  FTP files: 0, per second: 0, peak per second: 0
  FTP failure: 0
Export channel: 0
  Current server: Primary
  Primary server state: Unknown, Secondary server state: OK
Export channel: 1

```

Current server: Primary
 Primary server state: Unknown, Secondary server state: OK

show services flow-collector interface all terse

```
user@host> show services flow-collector interface all terse
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
  Packets      Bytes      Flows Uncompressed   Compressed   FTP bytes  FTP files
                Bytes      Bytes
    4384    6659616    131070    13742307    3786177    3786247      1

Flow collector interface: cp-6/3/0
Interface state: Collecting flows
  Packets      Bytes      Flows Uncompressed   Compressed   FTP bytes  FTP files
                Bytes      Bytes
      0         0         0         0         0         70         0
```

show services flow-collector interface extensive

```
user@host> show services flow-collector interface cp-5/2/0 extensive
Flow collector interface: cp-5/2/0
Interface state: Collecting flows
Memory:
  Used: 458311860, Free: 40810008
Input:
  Packets: 922629, per second: 2069, peak per second: 3266
  Bytes: 1376559252, per second: 3096940, peak per second: 4880051
  Flow records processed: 25764957, per second: 42564, peak per second: 98124
Allocation:
  Blocks allocated: 20862, per second: 31, peak per second: 72
  Blocks freed: 17161, per second: 40, peak per second: 202
  Blocks unavailable: 58786, per second: 652, peak per second: 1120
Files:
  Files created: 52, per second: 0, peak per second: 0
  Files exported: 42, per second: 0, peak per second: 0
  Files destroyed: 42, per second: 0, peak per second: 0
Throughput:
  Uncompressed bytes: 2592070401, per second: 7297307,
  peak per second: 8630023
  Compressed bytes: 659600068, per second: 1858458, peak per second: 2198471
Packet drops:
  No memory: 58786, Not IP: 0
  Not IPv4: 0, Too small: 0
  Fragments: 0, ICMP: 0
  TCP: 0, Unknown: 0
  Not JUNOS flow: 0
File Transfer:
  FTP bytes: 585981447, per second: 1313320, peak per second: 4857798
  FTP files: 48, per second: 0, peak per second: 0
  FTP failure: 8
Export channel: 0
  Current server: Primary
  Primary server state: FTP error, Secondary server state: Not configured
Export channel: 1
  Current server: Primary
  Primary server state: OK, Secondary server state: Not configured
```

show services rpm active-servers

| | |
|---------------------------------|---|
| Syntax | show services rpm active-servers |
| Release Information | Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| Description | Display the protocols and corresponding ports for which a router or switch is configured as a real-time performance monitoring (RPM) server. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show services rpm active-servers on page 2249 |
| Output Fields | Table 146 on page 2249 lists the output fields for the show services rpm active-servers command. Output fields are listed in the approximate order in which they appear. |

Table 146: show services rpm active-servers Output Fields

| Field Name | Field Description |
|-----------------------------------|---|
| Protocol | Protocol configured on the receiving probe server. The protocol can be the User Datagram Protocol (UDP) or the Transmission Control Protocol (TCP). |
| Port | Port configured on the receiving probe server. |
| Destination interface name | Output interface name for the probes. |

Sample Output

show services rpm active-servers

```
user@host> show services rpm active-servers
  Protocol: TCP, Port: 50000, Destination interface name: lt-0/0/0.0
  Protocol: UDP, Port: 50001, Destination interface name: lt-0/0/0.0
```

show services rpm history-results

| | |
|---------------------------------|---|
| Syntax | <pre>show services rpm history-results <brief detail> <owner <i>owner</i>> <since <i>time</i>> <test <i>name</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.</p> |
| Description | Display standard information about the results of the last 50 probes for each real-time performance monitoring (RPM) instance. |
| Options | <p>none—Display the results of the last 50 probes for all RPM instances.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>owner <i>owner</i>—(Optional) Display information for the specified probe owner.</p> <p>since <i>time</i>—(Optional) Display information from the specified time. Specify time as <i>yyyy-mm-dd.hh:mm:ss</i>.</p> <p>test <i>name</i>—(Optional) Display information for the specified test.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show services rpm history-results on page 2251</p> <p>show services rpm history-results detail on page 2251</p> |
| Output Fields | Table 147 on page 2250 lists the output fields for the show services rpm history-results command. Output fields are listed in the approximate order in which they appear. |

Table 147: show services rpm history-results Output Fields

| Field Name | Field Description | Level of Output |
|------------------------|--|-----------------|
| Owner | Probe owner. | All levels |
| Test | Name of a test for a probe instance. | All levels |
| Probe received | Timestamp when the probe result was determined. | All levels |
| Round trip time | Average ping round-trip time (RTT), in microseconds. | All levels |
| Probe results | <p>Result of a particular probe performed by a remote host. The following information is contained in the results:</p> <ul style="list-style-type: none"> Response received—Timestamp when the probe result was determined. Rtt—Average ping round-trip time (RTT), in microseconds. | detail |

Table 147: show services rpm history-results Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------------------|--|-----------------|
| Results over current test | Displays the results for the current test by probe at the time each probe was completed, as well as the status of the current test at the time the probe was completed. | detail |
| Probes sent | Number of probes sent with the current test. | detail |
| Probes received | Number of probe responses received within the current test. | detail |
| Loss percentage | Percentage of lost probes for the current test. | detail |
| Measurement | <p>Increment of measurement. Possible values are round-trip time delay and, for the probe type icmp-pin-timestamp, the egress and ingress delay:</p> <ul style="list-style-type: none"> • Minimum—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. • Maximum—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. • Average—Average RTT, ingress delay, or egress delay measured over the course of the current test. • Jitter—Difference, in microseconds, between the maximum and minimum RTT measured over the course of the current test. • Stddev—Standard deviation of the round-trip time, in microseconds, measured over the course of the current test. | detail |

Sample Output

show services rpm history-results

```

user@host> show services rpm history-results
      Owner, Test                Probe received                Round trip time
p1, t1                Wed Aug 12 01:02:35 2009                315 usec
p1, t1                Wed Aug 12 01:02:36 2009                266 usec
p1, t1                Wed Aug 12 01:02:37 2009                314 usec
p1, t1                Wed Aug 12 01:02:38 2009                388 usec
p1, t1                Wed Aug 12 01:02:39 2009                316 usec
p1, t1                Wed Aug 12 01:02:40 2009                271 usec
p1, t1                Wed Aug 12 01:02:41 2009                314 usec
p1, t1                Wed Aug 12 01:02:42 2009                1180 usec

```

show services rpm history-results detail

```

user@host> show services rpm history-results detail
Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
Probe results:
  Response received, Wed Aug 12 01:02:35 2009,
  Client and server hardware timestamps
  Rtt: 315 usec
Results over current test:
  Probes sent: 1, Probes received: 1, Loss percentage: 0
Measurement: Round trip time
  Samples: 1, Minimum: 315 usec, Maximum: 315 usec, Average: 315 usec,
  Peak to peak: 0 usec, Stddev: 0 usec, Sum: 315 usec

```

Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
Probe results:
Response received, Wed Aug 12 01:02:36 2009,
Client and server hardware timestamps
Rtt: 266 usec, Round trip jitter: -50 usec,
Round trip interarrival jitter: 3 usec
Results over current test:
Probes sent: 2, Probes received: 2, Loss percentage: 0
Measurement: Round trip time
Samples: 2, Minimum: 266 usec, Maximum: 315 usec, Average: 291 usec,
Peak to peak: 49 usec, Stddev: 24 usec, Sum: 581 usec
Measurement: Negative round trip jitter
Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,
Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec

Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
Probe results:
Response received, Wed Aug 12 01:02:37 2009,
Client and server hardware timestamps
Rtt: 314 usec, Round trip jitter: 49 usec,
Round trip interarrival jitter: 6 usec
Results over current test:
Probes sent: 3, Probes received: 3, Loss percentage: 0
Measurement: Round trip time
Samples: 3, Minimum: 266 usec, Maximum: 315 usec, Average: 298 usec,
Peak to peak: 49 usec, Stddev: 23 usec, Sum: 895 usec
Measurement: Positive round trip jitter
Samples: 1, Minimum: 49 usec, Maximum: 49 usec, Average: 49 usec,
Peak to peak: 0 usec, Stddev: 0 usec, Sum: 49 usec
Measurement: Negative round trip jitter
Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,
Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec

Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
Probe results:
Response received, Wed Aug 12 01:02:38 2009,
Client and server hardware timestamps
Rtt: 388 usec, Round trip jitter: 74 usec,
Round trip interarrival jitter: 10 usec
Results over current test:
Probes sent: 4, Probes received: 4, Loss percentage: 0
Measurement: Round trip time
Samples: 4, Minimum: 266 usec, Maximum: 388 usec, Average: 321 usec,
Peak to peak: 122 usec, Stddev: 44 usec, Sum: 1283 usec
Measurement: Positive round trip jitter
Samples: 2, Minimum: 49 usec, Maximum: 74 usec, Average: 62 usec,
Peak to peak: 25 usec, Stddev: 12 usec, Sum: 123 usec
Measurement: Negative round trip jitter
Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,
Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec

show services rpm probe-results

| | |
|---------------------------------|--|
| Syntax | show services rpm probe-results
<owner <i>owner</i> >
<test <i>name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 13.2 for PTX Series Packet Transport Series Routers. |
| Description | Display the results of the most recent real-time performance monitoring (RPM) probes. |
| Options | none —Display all results of the most recent RPM probes.

owner <i>owner</i> —(Optional) Display information for the specified probe owner.

test <i>name</i> —(Optional) Display information for the specified test. |
| Required Privilege Level | view |
| List of Sample Output | show services rpm probe-results on page 2256
show services rpm probe-results (BGP Neighbor Discovery) on page 2258 |
| Output Fields | Table 148 on page 2253 lists the output fields for the show services rpm probe-results command. Output fields are listed in the approximate order in which they appear. |

Table 148: show services rpm probe-results Output Fields

| Field Name | Field Description |
|-----------------------|---|
| Owner | Owner name. When you configure the probe owner statement at the [edit services rpm] hierarchy level, this field displays the configured owner name. When you configure BGP neighbor discovery through RPM, the output for this field is Rpm-Bgp-Owner . |
| Test | Name of a test representing a collection of probes. When you configure the test test-name statement at the [edit services rpm probe owner] hierarchy level, the field displays the configured test name. When you configure BGP neighbor discovery through RPM, the output for this field is Rpm-BGP-Test-<i>n</i> , where <i>n</i> is a cumulative number. |
| Target address | Destination address used for the probes. |
| Source address | Source address used for the probes. |
| Probe type | Protocol configured on the receiving probe server: http-get , http-metadata-get , icmp-ping , icmp-ping-timestamp , tcp-ping , udp-ping , or udp-ping-timestamp . |
| Test size | Number of probes within a test. |

Table 148: show services rpm probe-results Output Fields (*continued*)

| Field Name | Field Description |
|----------------------------------|---|
| Routing Instance Name | <p>(BGP neighbor discovery) Name of the configured (if any) routing instance, logical system name, or both, in which the probe is configured:</p> <ul style="list-style-type: none"> When a routing instance is defined within a logical system, the logical system name is followed by the routing instance name. A slash (/) is used to separate the two entities. For example, if the routing instance called R1 is configured within the logical system called LS, the name in the output field is LS/R1. When a routing instance is configured but the default logical system is used, the name in the output field is the name of the routing instance. When a logical system is configured but the default routing instance is used, the name in the output field is the name of the logical system followed by default. A slash (/) is used to separate the two entities. For example, LS/default. |
| Probe results | <p>Raw measurement of a particular probe sample done by a remote host. This data is provided separately from the calculated results. The following information is contained in the raw measurement:</p> <ul style="list-style-type: none"> Response received—Timestamp when the probe result was determined. Client and server hardware timestamps—If timestamps are configured, an entry appears at this point. Rtt—Average ping round-trip time (RTT), in microseconds. Egress jitter—Egress jitter, in microseconds. Ingress jitter—Ingress jitter, in microseconds. Round trip jitter—Round-trip jitter, in microseconds. Egress interarrival jitter—Egress interarrival jitter, in microseconds. Ingress interarrival jitter—Ingress interarrival jitter, in microseconds. Round trip interarrival jitter—Round-trip interarrival jitter, in microseconds. |
| Results over current test | <p>Probes are grouped into tests, and the statistics are calculated for each test. If a test contains 10 probes, the average, minimum, and maximum results are calculated from the results of those 10 probes. If the command is issued while the test is in progress, the statistics use information from the completed probes.</p> <ul style="list-style-type: none"> Probes sent—Number of probes sent within the current test. Probes received—Number of probe responses received within the current test. Loss percentage—Percentage of lost probes for the current test. Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type icmp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> Samples—Number of probes. Minimum—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. Maximum—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. Average—Average RTT, ingress delay, or egress delay measured over the course of the current test. Peak to peak—Peak-to-peak difference, in microseconds. Stddev—Standard deviation, in microseconds. Sum—Statistical sum. |

Table 148: show services rpm probe-results Output Fields (*continued*)

| Field Name | Field Description |
|-------------------------------|--|
| Results over last test | <p>Results for the most recently completed test. If the command is issued while the first test is in progress, this information is not displayed</p> <ul style="list-style-type: none"> • Probes sent—Number of probes sent for the most recently completed test. • Probes received—Number of probe responses received for the most recently completed test. • Loss percentage—Percentage of lost probes for the most recently completed test. • Test completed—Time the most recent test was completed. • Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type icmp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> • Samples—Number of probes. • Minimum—Minimum RTT, ingress delay, or egress delay measured for the most recently completed test. • Maximum—Maximum RTT, ingress delay, or egress delay measured for the most recently completed test. • Average—Average RTT, ingress delay, or egress delay measured for the most recently completed test. • Peak to peak—Peak-to-peak difference, in microseconds. • Stddev—Standard deviation, in microseconds. • Sum—Statistical sum. |
| Results over all tests | <p>Displays statistics made for all the probes, independently of the grouping into tests, as well as statistics for the current test.</p> <ul style="list-style-type: none"> • Probes sent—Number of probes sent in all tests. • Probes received—Number of probe responses received in all tests. • Loss percentage—Percentage of lost probes in all tests. • Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe types icmp-ping-timestamp and udp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> • Samples—Number of probes. • Minimum—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. • Maximum—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. • Average—Average RTT, ingress delay, or egress delay measured over the course of the current test. • Peak to peak—Peak-to-peak difference, in microseconds. • Stddev—Standard deviation, in microseconds. • Sum—Statistical sum. |

Table 148: show services rpm probe-results Output Fields (*continued*)

| Field Name | Field Description |
|--------------------|---|
| Error Stats | <p>Displays error statistics for each probe.</p> <ul style="list-style-type: none"> • Invalid client rcv timestamp—Number of client receive timestamp less than client send timestamp. • Invalid server send timestamp—Number of server send timestamp less than server receive timestamp. • Invalid server processing time—Number of server side spent time greater than RTT. <p>NOTE: Error Stats is displayed in the output only if non-zero statistics exists.</p> |

Sample Output

show services rpm probe-results

```

user@host> show services rpm probe-results
Owner: ADSN-J4300.ADSN-J2300.D2, Test: 75300002
Target address: 172.16.54.172, Source address: 10.206.0.1,
Probe type: udp-ping-timestamp, Test size: 10 probes
Probe results:
  Response received, Tue Feb  6 14:53:15 2007,
  Client and server hardware timestamps
  Rtt: 575 usec, Egress jitter: 5 usec, Ingress jitter: 8 usec,
  Round trip jitter: 12 usec, Egress interarrival jitter: 8 usec,
  Ingress interarrival jitter: 7 usec, Round trip interarrival jitter: 7 usec,

  Round trip interarrival jitter: 669 usec
Results over current test:
  Probes sent: 10, Probes received: 10, Loss percentage: 0
  Measurement: Round trip time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
  Measurement: Positive round trip jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
  Measurement: Negative round trip jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
  Measurement: Egress time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
  Measurement: Positive Egress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
  Measurement: Negative Egress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
  Measurement: Ingress time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
  Measurement: Positive Ingress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
  Measurement: Negative Ingress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Results over last test:
  Probes sent: 10, Probes received: 10, Loss percentage: 0

```

```

Test completed on Tue Feb 6 14:53:16 2007
Measurement: Round trip time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive round trip jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative round trip jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Measurement: Egress time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Egress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Egress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Measurement: Ingress time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Ingress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Ingress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Results over all tests:
Probes sent: 560, Probes received: 560, Loss percentage: 0
Measurement: Round trip time
  Samples: 560, Minimum: 805 usec, Maximum: 3114 usec, Average: 1756 usec,

  Peak to peak: 2309 usec, Stddev: 519 usec, Sum: xxxx usec
Measurement: Positive round trip jitter
  Samples: 257, Minimum: 0 usec, Maximum: 2054 usec, Average: 597 usec,
  Peak to peak: 2054 usec, Stddev: 427 usec, Sum: xxxx usec
Measurement: Negative round trip jitter
  Samples: 302, Minimum: 1 usec, Maximum: 1812 usec, Average: 511 usec,
  Peak to peak: 1811 usec, Stddev: 408 usec, Sum: xxxx usec
Measurement: Egress time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Egress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Egress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Measurement: Ingress time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Ingress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Ingress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Error Stats:
  Invalid client rcv timestamp: 3, Invalid server send timestamp: 0
  Invalid server processing time: 0

```

show services rpm probe-results (BGP Neighbor Discovery)

```
user@host> show services rpm probe-results
Owner: Rpm-Bgp-Owner, Test: Rpm-Bgp-Test-1
Target address: 10.209.152.37, Probe type: icmp-ping, Test size: 5 probes
Routing Instance Name: LS1/RI1
Probe results:
  Response received, Fri Oct 28 05:20:23 2005
  Rtt: 662 usec
Results over current test:
  Probes sent: 5, Probes received: 5, Loss percentage: 0
  Measurement: Round trip time
    Minimum: 529 usec, Maximum: 662 usec, Average: 585 usec,
    Jitter: 133 usec, Stddev: 53 usec
Results over all tests:
  Probes sent: 5, Probes received: 5, Loss percentage: 0
  Measurement: Round trip time
    Minimum: 529 usec, Maximum: 662 usec, Average: 585 usec,
    Jitter: 133 usec, Stddev: 53 usec
```

show services rpm rfc2544-benchmarking

| | |
|---------------------------------|---|
| Syntax | <pre>show services rpm rfc2544-benchmarking <aborted-tests (test-id test-id brief detail)> <active-tests (test-id test-id brief detail)> <completed-tests (test-id test-id brief detail)> <summary></pre> |
| Release Information | <p>Command introduced in Junos OS Release 12.3X52 for ACX Series routers.</p> <p>Command introduced in Junos OS Release 13.3R1 for MX104 3D Universal Edge Routers.</p> |
| Description | <p>Display information about the results of each category or state of the RFC 2544-based benchmarking test, such as aborted tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance. You can view the results of each test state for all of the configured test IDs or for a specific test ID. Also, you can display statistics about the total number of tests of each state for a high-level, quick analysis. The values in the output displayed vary, depending on the state in which the test is passing through, when you issue the command.</p> <p>You can view the test results of multiple test IDs at the same time by entering the IDs in a single command. If you enter multiple test ID values, you must separate each number with a space.</p> |
| Options | <p>aborted-tests—Display the list of tests that were aborted or stopped. This list includes tests that failed due to various error conditions and tests that you terminated by entering the test service rpm rfc2544-benchmarking test test-name stop command. The Status field in the output specifies the reason for the termination of the test.</p> <p>test-id test-id—Unique identifier of the test for which the test results must be displayed.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>active-tests—Display the results of the set of tests that are currently running.</p> <p>completed-tests—Display the results of the set of tests that were successfully completed. A completed test is one that passes through all the test steps or states specified in RFC 2544. A test that is marked as completed after it went through all the states from the beginning to the end can still be reported as a failed test. For example, a failed test can be a test that sends the desired number of packets, but does not receive the frames back from the other end.</p> <p>summary—(Optional) Display summary output.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • Configuring an RFC 2544-Based Benchmarking Test on page 989 • RFC2544-Based Benchmarking Tests Overview on page 983 • rfc2544-benchmarking on page 1740 |

List of Sample Output [show services rpm rfc2544-benchmarking summary on page 2261](#)
[show services rpm rfc2544-benchmarking aborted-tests \(ACX Series router\) on page 2261](#)
[show services rpm rfc2544-benchmarking completed-tests \(ACX Series router\) on page 2261](#)
[show services rpm rfc2544-benchmarking active-tests \(ACX Series router\) on page 2262](#)
[show services rpm rfc2544-benchmarking aborted-tests \(MX104 router\) on page 2262](#)
[show services rpm rfc2544-benchmarking completed-tests \(MX104 router\) on page 2262](#)
[show services rpm rfc2544-benchmarking active-tests \(MX104 router\) on page 2263](#)

Output Fields [Table 149 on page 2260](#) lists the output fields for the **show services rpm rfc2544-benchmarking (aborted-tests | active-tests | completed-tests)** command. Output fields are listed in the approximate order in which they appear.

Table 149: show services rpm rfc2544-benchmarking Output Fields

| Field Name | Field Description |
|-------------------------|--|
| Test information | Details of the performed RFC 2544 benchmarking test. |
| Test id | Unique identifier configured for the test. |
| Test name | Name configured for the test. |
| Test type | The type of statistical detail that is collected for the test, based on the configured test type. Throughput-related, latency, frame-loss, or back-to-back frames-related information is displayed for ACX Series routers. Reflected packets-related information is displayed for MX104 routers.. |
| Test mode | Mode configured for the test on the router. Test modes are: <ul style="list-style-type: none"> Initiate-and-Terminate: Test frames are initiated from one end and terminated at the same end. This mode requires a reflector to be configured at the peer end to enable the test frames to be returned to the source. This mode is supported only on ACX Series routers Reflect: Test frames that originate from one end are reflected at the other end on the selected service, such as IPv4 or Ethernet. |
| Test packet size | Size of the test packets in bytes. This field is valid only when the test mode is Initiate-and-Terminate. |
| Test state | State of the test that is in progress or active when the output is displayed. |
| Status | Indicates whether the test is currently in progress or has been terminated. This field is displayed for tests that are in progress or were aborted by entering the test services rpm rfc2544-benchmarking test <test-name test-id> stop command. |
| Test start time | Time at which the test started in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS). |
| Test finish time | Time at which the test completed. |

Table 149: show services rpm rfc2544-benchmarking Output Fields (*continued*)

| Field Name | Field Description |
|---------------------------|---|
| Counters last cleared | Date, time, and how long ago the statistics for the test were cleared. The format is <i>year-month-day hour:minute:second:timezone</i> (<i>hour:minute:second</i> ago). For example, 2010-05-17 07:51:28 PDT (00:04:33 ago). If you did not clear the statistics previously at any point, Never is displayed. |
| Number of active tests | Total number of tests that are currently running. |
| Number of completed tests | Total number of tests that were successfully completed |
| Number of aborted tests | Total number of tests that were aborted or halted. |

Sample Output

show services rpm rfc2544-benchmarking summary

```
user@host> show services rpm rfc2544-benchmarking summary
```

```
Rfc2544 tests summary :
```

```
    Number of active tests: 0, Number of completed tests: 4, Number of aborted tests: 52
```

This output indicates that no test iteration is currently in progress (at the time of issue of the command), 4 tests were completed successfully, and 52 tests were halted.

show services rpm rfc2544-benchmarking aborted-tests (ACX Series router)

```
user@host> show services rpm rfc2544-benchmarking aborted-tests
```

```
Test information :
```

```
    Test id: 1, Test name: test1, Test type: Throughput
    Test mode: Initiate-and-Terminate
    Test packet size: 64 1280
    Test state: RFC2544_TEST_STATE_STOPPED
    Status: User-aborted-via-cli
    Test start time: 2005-08-05 03:19:58 UTC
    Test finish time: 2005-08-05 03:20:00 UTC
    Counters last cleared: Never
```

```
    Test id: 2, Test name: test1, Test type: Throughput
    Test mode: Initiate-and-Terminate
    Test packet size: 64 1280
    Test state: RFC2544_TEST_STATE_STOPPED
    Status: User-aborted-via-cli
    Test start time: 2005-08-05 03:20:00 UTC
    Test finish time: 2005-08-05 03:20:02 UTC
    Counters last cleared: Never
```

show services rpm rfc2544-benchmarking completed-tests (ACX Series router)

```
user@host> show services rpm rfc2544-benchmarking completed-tests
```

```
Test information :
  Test id: 18, Test name: test1, Test type: Throughput
  Test mode: Initiate-and-Terminate
  Test packet size: 64 1280
  Test state: RFC2544_TEST_STATE_COMPLETED
  Test start time: 2005-08-05 03:20:34 UTC
  Test finish time: 2005-08-05 03:21:23 UTC
  Counters last cleared: Never
```

show services rpm rfc2544-benchmarking active-tests (ACX Series router)

```
user@host> show services rpm rfc2544-benchmarking active-tests
Test information :
  Test id: 57, Test name: test1, Test type: Back-Back-Frames
  Test mode: Initiate-and-Terminate
  Test packet size: 64 1280
  Test state: RFC2544_TEST_STATE_RUNNING
  Status: Running
  Test start time: 2005-08-05 20:15:41 UTC
  Test finish time: TEST_RUNNING
  Counters last cleared: Never
```

show services rpm rfc2544-benchmarking aborted-tests (MX104 router)

```
user@host> show services rpm rfc2544-benchmarking aborted-tests
Test information :
  Test id: 1, Test name: prof_tput1, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: TEST_STATE_STOPPED
  Status: Test-intf-ifl-change
  Test start time: 2013-12-16 22:54:27 PST
  Test finish time: 2013-12-16 23:30:28 PST
  Counters last cleared: Never

  Test id: 2, Test name: prof_tput1, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: TEST_STATE_STOPPED
  Status: User-aborted-via-cli
  Test start time: 2013-12-16 23:31:06 PST
  Test finish time: 2013-12-16 23:36:22 PST
  Counters last cleared: Never

  Test id: 3, Test name: prof_tput1, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: TEST_STATE_STOPPED
  Status: User-aborted-via-cli
  Test start time: 2013-12-16 23:36:24 PST
  Test finish time: 2013-12-17 01:49:24 PST
  Counters last cleared: Never
```

show services rpm rfc2544-benchmarking completed-tests (MX104 router)

```
user@host> show services rpm rfc2544-benchmarking completed-tests
Test information :
  Test id: 18, Test name: test1, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: TEST_STATE_COMPLETED
```


Test start time: 2005-08-05 03:20:34 UTC
Test finish time: 2005-08-05 03:21:23 UTC
Counters last cleared: Never

show services rpm rfc2544-benchmarking active-tests (MX104 router)

```
user@host> show services rpm rfc2544-benchmarking active-tests
Test information :
  Test id: 4, Test name: prof_tput1, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: TEST_STATE_RUNNING
  Status: Running
  Test start time: 2013-12-17 01:49:26 PST
  Test finish time: TEST_RUNNING
  Counters last cleared: Never
```

show services rpm rfc2544-benchmarking test-id

| | |
|---------------------------------|---|
| Syntax | <code>show services rpm rfc2544-benchmarking test-id <i>test-id</i></code>
<code><brief detail></code> |
| Release Information | Command introduced in Junos OS Release 12.3X52 for ACX Series routers.
Command introduced in Junos OS Release 13.3R1 for MX104 3D Universal Edge Routers. |
| Description | Display information about the results of the RFC 2544-based benchmarking test for a specific test ID for each real-time performance monitoring (RPM) instance. The values in the output displayed vary, depending on the state in which the test is passing through, when you issue the command. |
| Options | none —Display brief information about a specific test ID of the benchmarking test.

test-id <i>test-id</i> —Unique identifier of the test for which the test results must be displayed.

brief detail —(Optional) Display the specified level of output. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• Configuring an RFC 2544-Based Benchmarking Test on page 989• RFC2544-Based Benchmarking Tests Overview on page 983• rfc2544-benchmarking on page 1740 |
| List of Sample Output | show services rpm rfc2544-benchmarking test-id detail (Throughput Test on ACX Series routers) on page 2272
show services rpm rfc2544-benchmarking test-id detail (Latency Test on ACX Series routers) on page 2273
show services rpm rfc2544-benchmarking test-id detail (Frame Loss Test on ACX Series routers) on page 2276
show services rpm rfc2544-benchmarking test-id detail (Back-to-Back Frames Test on ACX Series routers) on page 2277
show services rpm rfc2544-benchmarking test-id detail (Reflection Test on MX104 routers) on page 2278
show services rpm rfc2544-benchmarking test-id brief (Reflection Test on MX104 routers) on page 2279
show services rpm rfc2544-benchmarking test-id detail (Reflection Test on MX104 routers) on page 2279
show services rpm rfc2544-benchmarking test-id brief (Reflection Test on MX104 routers) on page 2280 |
| Output Fields | Table 150 on page 2265 lists the output fields for the show services rpm rfc2544-benchmarking test-id command. Output fields are listed in the approximate order in which they appear. |

Table 150: show services rpm rfc2544-benchmarking test-id Output Fields

| Field Name | Field Description | Level of Output |
|-----------------------------------|---|-----------------|
| Test information | Details of the performed RFC 2544 benchmarking test. | None specified |
| Test id | Unique identifier configured for the test. | None specified |
| Test name | Name configured for the test. | None specified |
| Test type | The type of statistical detail that is collected for the test, based on the configured test type. Throughput-related, latency, frame-loss, or back-to-back frames-related information is displayed for ACX Series routers. Reflected packets-related information is displayed for MX104 routers. | None specified |
| Test mode | Mode configured for the test on the router. Test modes are: <ul style="list-style-type: none"> Initiate-and-Terminate: Test frames are initiated from one end and terminated at the same end. This mode requires a reflector to be configured at the peer end to enable the test frames to be returned to the source. This mode is supported only on ACX Series routers. Reflect: Test frames that originate from one end are reflected at the other end on the selected service, such as IPv4 or Ethernet. | None specified |
| Test packet size | Size of the test packets in bytes. This field is valid only when the test mode is Initiate-and-Terminate. | None specified |
| Test state | State of the test that is in progress or active when the output is displayed. For details about the states, see <i>RFC 2544-Based Benchmarking Test States</i> . | None specified |
| Status | Indicates whether the test is currently in progress or has been terminated. | None specified |
| Test start time | Time at which the test started in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS). | None specified |
| Test finish time | Time at which the test completed. | None specified |
| Counters last cleared | Date, time, and how long ago the statistics for the test were cleared. The format is <i>year-month-day hour:minute:second:timezone (hour:minute:second ago)</i> . For example, 2010-05-17 07:51:28 PDT (00:04:33 ago). If you did not clear the statistics previously at any point, Never is displayed. | None specified |
| Test-profile Configuration | (ACX Series routers only) Details of the specified test profile | detail |
| Test-profile name | (ACX Series routers only) Name of the configured test profile that contains the parameters for the test | detail |
| Test packet size | (ACX Series routers only) Size of the test packets in bytes | detail |
| Theoretical max bandwidth | (ACX Series routers only) Theoretical maximum bandwidth configured for the test. This value is typically set to the bandwidth of the server being tested. Valid values are 1 Kbps through 1,000,000 Kbps (1 Gbps). The value defined is the highest bandwidth value tested for this test. | detail |

Table 150: show services rpm rfc2544-benchmarking test-id Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------------------|--|-----------------|
| Test Configuration | Details of the configured test ID. | detail |
| Test mode | Mode configured for the test. Test modes are Initiate-and-Terminate and Reflect. | detail |
| Duration in seconds | Period in seconds for which the test has been performed. | detail |
| Test family | The underlying service on which the test is run. Test families are: <ul style="list-style-type: none"> • INET: Indicates that the test is run on a IPV4 service. • CCC: Indicates that the test is run on a circuit cross-connect (CCC) or pseudowire service. | detail |
| Routing Instance Name | (ACX Series routers only) Name of the routing instance for the test | detail |
| Inet family Configuration | Details of the configured inet family for an IPv4 service | detail |
| Egress Interface | Name of the egress interface from which the test frames are sent | detail |
| Source ipv4 address | Source IPv4 address used in the IP header of the generated test frame. | detail |
| Destination ipv4 address | Destination IPv4 address used in the IP header of the generated test frame. | detail |
| Source udp port | Source UDP port number used in the UDP header of the generated test frame. | detail |
| Destination udp port | Destination UDP port number used in the UDP header of the generated test frame. | detail |
| Ccc family Configuration | Details of the configured CCC family for an Ethernet service | detail |
| Source MAC address | (ACX Series routers only) Source MAC address used in generated test frames for a CCC or Ethernet pseudowire service. | detail |
| Destination MAC address | (ACX Series routers only) Destination MAC address used in generated test frames for a CCC or Ethernet pseudowire service. | detail |
| Ivlan-id | (ACX Series routers only) Inner VLAN ID for test-frames. | detail |
| Ovlan-id | (ACX Series routers only) Outer VLAN ID for test-frames. | detail |
| Direction egress | Test is run in the egress direction of the interface (NNI) | detail |
| Direction ingress | Test is run in the ingress direction of the interface (UNI) | detail |

Table 150: show services rpm rfc2544-benchmarking test-id Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|--|-----------------|
| Rfc2544 throughput test information | (ACX Series routers only) Details of the throughput test | detail |
| Initial test load percentage | Percentage of the steady state load for the test. | detail |
| Test iteration mode | Mode of the test iteration: Binary or step-down. | detail |
| Test iteration step percent | The test step percentage for tests. If not specified, the default step-percent is 10 percent. This parameter is ignored for all type of tests other than frame-loss tests. | detail |
| Theoretical max bandwidth | The theoretical limit of the media for the frame size configured for the test. This value is typically set to the bandwidth of the server being tested. | detail |
| Test packet size: | Packet size of the test frames in bytes. | detail |
| Iteration | Number of the test iteration. | detail |
| Duration (sec) | Period in seconds for which the test iteration is run | detail |
| Elapsed time | Amount of time that has passed, in seconds, since the start of the test. | detail |
| pps | Total count of packets-per-second (pps) transmitted during the test. | detail |
| Tx Packets | Number of transmitted test packets. | detail |
| Rx Packets | Number of received test packets. | detail |
| Tx Bytes | Number of transmitted bytes. | detail |
| Rx Bytes | Number of received bytes. | detail |
| Percentage throughput | Percentage of throughput for the test iteration. | detail |
| Result of the iteration runs (Throughput) : | Results of the completed throughput test for a particular packet size. | detail |
| Best iteration | Number of the iteration with the highest throughout, among the listed iterations. | detail |
| Best iteration (pps) | Packets-per-second (pps) count of the iteration with the highest throughout, among the listed iterations. | detail |
| Best iteration throughput | Percentage of throughput of the iteration with the highest throughout, among the listed iterations. | detail |

Table 150: show services rpm rfc2544-benchmarking test-id Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|---|-----------------------|
| RFC2544 Throughput test results summary | Consolidated information of the throughput test. | detail summary |
| Packet Size | Size of the test packet in bytes. | detail summary |
| Theoretical rate (pps) | Theoretical frame rate in packets-per-second. | detail summary |
| Tx Packets | Number of transmitted packets. | detail summary |
| Rx Packets | Number of received packets. | detail summary |
| Offered throughput (percentage) | The offered throughput in percentage of the chosen service (such as Layer 3 or Ethernet pseudowire). | detail summary |
| Measured bandwidth (kbps) | Available bandwidth of the service based on the calculated throughput. | detail summary |
| Rfc2544 latency test information : | (ACX Series routers only) Details of the latency test | detail |
| Theoretical max bandwidth | Theoretical maximum bandwidth configured for the test. This value is typically set to the bandwidth of the server being tested. Valid values are 1 Kbps through 1,000,000 Kbps (1 Gbps). The value defined is the highest bandwidth value used for this test. | detail |
| Initial test load percentage | Percentage of the steady state load for the test. | detail |
| Duration in seconds | Period in seconds for which the test has been performed. | detail |
| Test packet size | Size of the test packet in bytes. | detail |
| Iteration | Number of the test iteration. | detail |
| Duration (sec) | Period in seconds for which the test iteration is run. | detail |
| Elapsed time | Amount of time that has passed, in seconds, since the start of the test. | detail |
| pps | Total count of packets-per-second (pps) transmitted during the test. | detail |
| Tx Packets | Number of transmitted test packets. | detail |
| Rx Packets | Number of received test packets. | detail |
| Latency | Displays the latency parameters. | detail |

Table 150: show services rpm rfc2544-benchmarking test-id Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|---|-----------------------|
| Min(ns) | Aggregated minimum latency in nanoseconds. | detail |
| Avg(ns) | Aggregated average latency in nanoseconds. | detail |
| Max(ns) | Aggregated maximum latency in nanoseconds. | detail |
| Probe(ns) | Aggregated probe latency in nanoseconds. | detail |
| Result of the iteration runs (Latency) | Results of the latency test completed for a particular packet size. | detail |
| Avg (min) Latency | Average of the minimum latency in nanoseconds. | detail |
| Avg (avg) latency | Average of the average latency in nanoseconds. | detail |
| Avg (Max) latency | Average of the maximum latency in nanoseconds. | detail |
| Avg (probe) latency | Average of the probe latency in nanoseconds. | detail |
| RFC2544 Latency test results summary: | Consolidated statistics of the latency test. | detail summary |
| Packet Size | Size of the test packet in bytes. | detail summary |
| Theoretical rate (pps) | Theoretical frame rate in packets-per-second. | detail summary |
| Tx Packets | Number of transmitted packets. | detail summary |
| Rx Packets | Number of received packets. | detail summary |
| Latency | Displays the latency parameters. | detail summary |
| Min(ns) | Aggregated minimum latency in nanoseconds. | detail summary |
| Avg(ns) | Aggregated average latency in nanoseconds. | detail summary |
| Max(ns) | Aggregated maximum latency in nanoseconds. | detail summary |
| Probe(ns) | Aggregated probe latency in nanoseconds. | detail summary |
| Rfc2544 Back-Back test information : | (ACX Series routers only) Details of the back-to-back frames or bursty frames test. | detail |
| Initial burst length: | Length of the first burst when test frames are sent, as a measure of number of seconds at the rate of Kbps. | detail |

Table 150: show services rpm rfc2544-benchmarking test-id Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|--|-----------------------|
| Test iteration mode : | Mode of the test iteration: Binary or step-down. | detail |
| Test iteration step percent | The test step percentage for tests. If not specified, the default step-percent is 10 percent. This parameter is ignored for all type of tests other than frame-loss tests. | detail |
| Theoretical max bandwidth | The theoretical limit of the media for the frame size configured for the test. This value is typically set to the bandwidth of the server being tested. | detail |
| Test packet size: | Packet size of the test frames in bytes. | detail |
| Iteration | Number of the test iteration. | detail |
| Burst Length (Packets) | Number of packets in the burst. | detail |
| Elapsed time | Amount of time that has passed, in seconds, since the start of the test. | detail |
| Tx Packets | Number of transmitted test packets. | detail |
| Rx Packets | Number of received test packets. | detail |
| Tx Bytes | Number of transmitted bytes. | detail |
| Rx Bytes | Number of received bytes. | detail |
| Result of the iteration runs : | Results of the back-to-back frames test completed for a certain packet size. | detail |
| Best iteration : | Number of the iteration with the longest burst. | detail |
| Measured burst (num sec) | Time in seconds of the burst of the iteration with the longest burst. | detail |
| Measured burst (num pkts) | Number of packets during the burst of the iteration with the longest burst. | detail |
| RFC2544 Back-Back test results summary: | Consolidated statistics of the back-to-back frames test. | detail summary |
| Packet Size | Size of the test packets in bytes. | detail summary |
| Measure Burst length (Packets) | Computed burst length in terms of number of packets. | detail summary |

Table 150: show services rpm rfc2544-benchmarking test-id Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|--|-----------------|
| Rfc2544 frame-loss test information : | (ACX Series routers only) Details of the frame-loss test. | detail |
| Initial burst length: | Length of the first burst when test frames are sent, as a measure of number of seconds at the rate of Kbps. | detail |
| Test iteration mode : | Mode of the test iteration: Binary or step-down. | detail |
| Test iteration step percent | The test step percentage for tests. If not specified, the default step-percent is 10 percent. This parameter is ignored for all type of tests other than frame-loss tests. | detail |
| Theoretical max bandwidth | The theoretical limit of the media for the frame size configured for the test. This value is typically set to the bandwidth of the server being tested. | detail |
| Test packet size | Size of the test packets in bytes. | detail |
| Iteration | Number of the test iteration. | detail |
| Duration (sec) | Period, in seconds, for which the test iteration is run. | detail |
| Offered throughput (percentage) | The offered throughput in percentage of the chosen service (such as Layer 3 or Ethernet pseudowire) | detail |
| Elapsed time | Amount of time that has passed, in seconds, since the start of the test. | detail |
| pps | Theoretical frame rate in packets-per-second. | detail |
| Tx Packets | Number of transmitted test packets. | detail |
| Rx Packets | Number of received test packets. | detail |
| Tx Bytes | Number of transmitted bytes. | detail |
| Rx Bytes | Number of received bytes. | detail |
| Frame-loss rate % | Percentage of frames that must been forwarded by the router under steady state (constant) load, but were not forwarded due to lack of resources. | detail |
| Result of the iteration runs : | Results of the frame-loss test completed for a certain packet size. | detail |
| Frame-loss rate (percent) : | Percentage of dropped frames for the specified packet size | detail |
| RFC2544 Frame-loss test results summary | Consolidated statistics of the frame-loss test | detail |

Table 150: show services rpm rfc2544-benchmarking test-id Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|--|-----------------|
| Packet Size | Size of the test packet in bytes. | detail summary |
| Theoretical rate (pps) | Theoretical frame rate in packets-per-second. | detail summary |
| Percentage throughput | Percentage of throughput for the test iteration. | detail summary |
| Tx Packets | Number of transmitted packets. | detail summary |
| Rx Packets | Number of received packets. | detail summary |
| Frame Loss rate percent | Percentage of dropped frames for the specified packet size | detail summary |

Sample Output

show services rpm rfc2544-benchmarking test-id detail (Throughput Test on ACX Series routers)

```

user@host> show services rpm rfc2544-benchmarking test-id 19 detail
Test information :
  Test id: 19, Test name: test1, Test type: Throughput
  Test mode: Initiate-and-Terminate
  Test packet size: 64 1280
  Test state: RFC2544_TEST_STATE_COMPLETED
  Test start time: 2005-07-29 10:25:00 UTC
  Test finish time: 2005-07-29 10:26:02 UTC
  Counters last cleared: Never

Test-profile Configuration:
  Test-profile name: prof_tput
  Test packet size: 64 1280
  Therotical max bandwidth : 993000 kbps

Test Configuration:
  Test mode: Initiate-and-Terminate
  Duration in seconds: 20
  Test family: INET
  Routing Instance Name: default

Inet family Configuration:
  Egress Interface : ge-0/1/1.0
  Source ipv4 address: 20.6.0.1
  Destination ipv4 address: 20.6.0.2
  Source udp port: 2020
  Destination udp port: 3030

Rfc2544 throughput test information :
  Initial test load percentage : 100.00 %
  Test iteration mode : Binary
  Test iteration step percent : 50.00 %
  Therotical max bandwidth : 993000 kbps

```

```

Test packet size: 64
Iteration Duration Elapsed pps      Tx      Rx      Tx      Rx
Percentage
(sec)   time      Packets  Packets  Bytes    Bytes
throughput
1       3       3       134918  404754  404754  27523272  27523272  10.00
%
2       20      20      1349184 26983501 26983501 1834878068 1834878068 100.00
%

```

Result of the iteration runs : Throughput Test complete for packet size 64
 Best iteration : 2, Best iteration (pps) : 1349184
 Best iteration throughput : 100.00 %

```

Test packet size: 1280
Iteration Duration Elapsed pps      Tx      Rx      Tx      Rx
Percentage
(sec)   time      Packets  Packets  Bytes    Bytes
throughput
1       3       3       9489   28467   28467   36551628  36551628  10.00
%
2       20      20      94896  1897920 1897920  2436929280 2436929280 100.00
%

```

Result of the iteration runs : Throughput Test complete for packet size 1280
 Best iteration : 2, Best iteration (pps) : 94896
 Best iteration throughput : 100.00 %

RFC2544 Throughput test results summary:

| Packet Size | Theoretical rate (pps) | Tx Packets | Rx Packets | Offered throughput (percentage) | Measured bandwidth (kbps) |
|-------------|------------------------|------------|------------|---------------------------------|---------------------------|
| 64 | 1349184 | 26983501 | 26983501 | 100.00 % | 993000 |
| 1280 | 94896 | 1897920 | 1897920 | 100.00 % | 993000 |

show services rpm rfc2544-benchmarking test-id detail (Latency Test on ACX Series routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 37 detail
```

Test information :

```

Test id: 37, Test name: test1, Test type: Latency
Test mode: Initiate-and-Terminate
Test packet size: 64 1280
Test state: RFC2544_TEST_STATE_COMPLETED
Test start time: 2005-07-29 10:26:41 UTC
Test finish time: 2005-07-29 10:36:15 UTC
Counters last cleared: Never

```

Test-profile Configuration:

```

Test-profile name: prof_latency
Test packet size: 64 1280
Theoretical max bandwidth : 993000 kbps

```

Test Configuration:

```

Test mode: Initiate-and-Terminate
Duration in seconds: 10
Test family: INET
Routing Instance Name: default

```

Inet family Configuration:

Egress Interface : ge-0/1/1.0
 Source ipv4 address: 20.6.0.1
 Destination ipv4 address: 20.6.0.2
 Source udp port: 2020
 Destination udp port: 3030

Rfc2544 latency test information :

Theoretical max bandwidth : 993000 kbps
 Initial test load percentage : 100.00 %
 Duration in seconds: 10

Test packet size: 64

| Iteration | Duration
(sec) | Elapsed
time | pps | Tx
Packets | Rx
Packets |
|-----------|-------------------|-----------------|---------|---------------|---------------|
| 1 | 3 | 3 | 134918 | 404754 | 404754 |
| 2 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 3 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 4 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 5 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 6 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 7 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 8 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 9 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 10 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 11 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 12 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 13 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 14 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 15 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 16 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 17 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 18 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 19 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 20 | 10 | 10 | 1349184 | 13491751 | 13491751 |
| 21 | 10 | 10 | 1349184 | 13491751 | 13491751 |

| ----- Latency ----- | | | |
|---------------------|---------|---------|-----------|
| Min(ns) | Avg(ns) | Max(ns) | Probe(ns) |
| 17464 | 18770 | 18880 | 18784 |
| 17472 | 18799 | 20488 | 18848 |
| 17472 | 18799 | 20416 | 18816 |
| 17472 | 18799 | 20440 | 18704 |
| 17464 | 18799 | 20376 | 18880 |
| 17464 | 18799 | 20232 | 18832 |
| 17464 | 18799 | 20400 | 18848 |
| 17472 | 18799 | 20240 | 18864 |
| 17472 | 18799 | 20264 | 18848 |
| 17464 | 18799 | 20264 | 18880 |
| 17472 | 18800 | 20320 | 18864 |
| 17464 | 18799 | 20176 | 18864 |
| 17464 | 18800 | 20248 | 18864 |
| 17464 | 18800 | 20272 | 18864 |
| 17464 | 18799 | 20472 | 18832 |
| 17464 | 18799 | 20256 | 18880 |
| 17464 | 18799 | 20336 | 18848 |
| 17464 | 18800 | 20688 | 18848 |
| 17472 | 18800 | 20504 | 18864 |
| 17464 | 18799 | 20448 | 18768 |
| 17472 | 18799 | 20240 | 18864 |

Result of the iteration runs : Latency Test complete for packet size 64

Avg (min) Latency : 17466
 Avg (avg) latency : 18799
 Avg (Max) latency : 20360
 Avg (probe) latency : 18844

Test packet size: 1280

| Iteration | Duration
(sec) | Elapsed
time | pps | Tx
Packets | Rx
Packets |
|-----------|-------------------|-----------------|-------|---------------|---------------|
| 1 | 3 | 3 | 9489 | 28467 | 28467 |
| 2 | 10 | 10 | 94896 | 948960 | 948960 |
| 3 | 10 | 10 | 94896 | 948960 | 948960 |
| 4 | 10 | 10 | 94896 | 948960 | 948960 |
| 5 | 10 | 10 | 94896 | 948960 | 948960 |
| 6 | 10 | 10 | 94896 | 948960 | 948960 |
| 7 | 10 | 10 | 94896 | 948960 | 948960 |
| 8 | 10 | 10 | 94896 | 948960 | 948960 |
| 9 | 10 | 10 | 94896 | 948960 | 948960 |
| 10 | 10 | 10 | 94896 | 948960 | 948960 |
| 11 | 10 | 10 | 94896 | 948960 | 948960 |
| 12 | 10 | 10 | 94896 | 948960 | 948960 |
| 13 | 10 | 10 | 94896 | 948960 | 948960 |
| 14 | 10 | 10 | 94896 | 948960 | 948960 |
| 15 | 10 | 10 | 94896 | 948960 | 948960 |
| 16 | 10 | 10 | 94896 | 948960 | 948960 |
| 17 | 10 | 10 | 94896 | 948960 | 948960 |
| 18 | 10 | 10 | 94896 | 948960 | 948960 |
| 19 | 10 | 10 | 94896 | 948960 | 948960 |
| 20 | 10 | 10 | 94896 | 948960 | 948960 |
| 21 | 10 | 10 | 94896 | 948960 | 948960 |

----- Latency -----

| Min(ns) | Avg(ns) | Max(ns) | Probe(ns) |
|---------|---------|---------|-----------|
| 68712 | 70031 | 70576 | 69456 |
| 68728 | 70344 | 71808 | 70512 |
| 68720 | 70344 | 71744 | 70352 |
| 68720 | 70344 | 71680 | 70112 |
| 68720 | 70345 | 71856 | 70352 |
| 68720 | 70344 | 71808 | 70384 |
| 68720 | 70344 | 71752 | 70480 |
| 68720 | 70344 | 71880 | 70112 |
| 68720 | 70344 | 71792 | 70320 |
| 68728 | 70345 | 73344 | 70336 |
| 68720 | 70344 | 71688 | 70560 |
| 68728 | 70345 | 71896 | 70496 |
| 68720 | 70344 | 71760 | 70096 |
| 68720 | 70344 | 71776 | 70320 |
| 68720 | 70344 | 71760 | 70400 |
| 68712 | 70345 | 71920 | 70352 |
| 68720 | 70344 | 71792 | 70576 |
| 68720 | 70345 | 71840 | 70320 |
| 68720 | 70344 | 71792 | 70368 |
| 68720 | 70345 | 71824 | 70464 |
| 68712 | 70345 | 71904 | 70512 |

Result of the iteration runs : Latency Test complete for packet size 1280

Avg (min) Latency : 68720
 Avg (avg) latency : 70344

```

Avg (Max) latency           : 71880
Avg (probe) latency         : 70371

```

RFC2544 Latency test results summary:

```

-----
Packet  Theoretical Tx      Rx      ----- Latency -----
Size   rate (pps)  Packets  Packets  Min(ns)  Avg(ns)  Max(ns)  Probe(ns)
64     1349184    269835020 269835020 17466    18799    20360    18844
1280   94896       18979200 18979200  68720    70344    71880    70371

```

show services rpm rfc2544-benchmarking test-id detail (Frame Loss Test on ACX Series routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 73 detail
```

Test information :

```

Test id: 73, Test name: test1, Test type: Frame-Loss
Test mode: Initiate-and-Terminate
Test packet size: 64 1280
Test state: RFC2544_TEST_STATE_COMPLETED
Test start time: 2005-07-29 10:38:41 UTC
Test finish time: 2005-07-29 10:41:19 UTC
Counters last cleared: Never

```

Test-profile Configuration:

```

Test-profile name: prof_fl
Test packet size: 64 1280
Theoretical max bandwidth : 993000 kbps

```

Test Configuration:

```

Test mode: Initiate-and-Terminate
Duration in seconds: 20
Test family: INET
Routing Instance Name: default

```

Inet family Configuration:

```

Egress Interface : ge-0/1/1.0
Source ipv4 address: 20.6.0.1
Destination ipv4 address: 20.6.0.2
Source udp port: 2020
Destination udp port: 3030

```

Rfc2544 frame-loss test information :

```

Initial test load percentage : 100.00 %
Test iteration mode : step-down
Test iteration step percent : 10 %
Theoretical max bandwidth : 993000 kbps

```

Test packet size: 64

| Iteration | Duration | Elapsed | Offered | pps | Tx | Rx | Tx | Rx |
|------------|----------|---------|-------------|---------|----------|----------|------------|-------|
| Frame-loss | | | | | | | | |
| | (sec) | time | throughput% | | Packets | Packets | Bytes | Bytes |
| | rate % | | | | | | | |
| 1 | 3 | 3 | 10.00 % | 134918 | 404754 | 404754 | 27523272 | |
| 27523272 | 0.00 % | | | | | | | |
| 2 | 20 | 20 | 100.00 % | 1349184 | 26983501 | 26983501 | 1834878068 | |
| 1834878068 | 0.00 % | | | | | | | |
| 3 | 20 | 20 | 100.00 % | 1349184 | 26983501 | 26983501 | 1834878068 | |
| 1834878068 | 0.00 % | | | | | | | |
| 4 | 20 | 20 | 100.00 % | 1349184 | 26983501 | 26983501 | 1834878068 | |
| 1834878068 | 0.00 % | | | | | | | |

Result of the iteration runs : Frame-loss test complete for packet size 64
 Frame-loss rate (percent) : 0.00 %

Test packet size: 1280

| Iteration | Duration
Frame-loss
(sec) | Elapsed
time | Offered
throughput% | pps | Tx
Packets | Rx
Packets | Tx
Bytes | Rx
Bytes |
|-----------|---------------------------------|-----------------|------------------------|-------|---------------|---------------|-------------|-------------|
| 1 | 3 | 3 | 10.00 % | 9489 | 404754 | 28467 | 36551628 | |
| 2 | 20 | 20 | 100.00 % | 94896 | 1897920 | 1897920 | 2436929280 | |
| 3 | 20 | 20 | 100.00 % | 94896 | 1897920 | 1897920 | 2436929280 | |
| 4 | 20 | 20 | 100.00 % | 94896 | 1897920 | 1897920 | 2436929280 | |

Result of the iteration runs : Frame-loss test complete for packet size 1280
 Frame-loss rate (percent) : 0.00 %

RFC2544 Frame-loss test results summary:

| Packet
Loss
Size
percent | Theoretical
rate (pps) | Percentage
throughput | Tx
Packets | Rx
Packets | Frame
rate |
|-----------------------------------|---------------------------|--------------------------|---------------|---------------|---------------|
| 64 | 1349184 | 100.00 % | 26983501 | 26983501 | 0.00 |
| 1280 | 94896 | 100.00 % | 1897920 | 1897920 | 0.00 |

show services rpm rfc2544-benchmarking test-id detail (Back-to-Back Frames Test on ACX Series routers)

```

user@host> show services rpm rfc2544-benchmarking test-id 55 detail
Test information :
  Test id: 55, Test name: test1, Test type: Back-Back-Frames
  Test mode: Initiate-and-Terminate
  Test packet size: 64 1280
  Test state: RFC2544_TEST_STATE_COMPLETED
  Test start time: 2005-07-29 10:36:54 UTC
  Test finish time: 2005-07-29 10:37:57 UTC
  Counters last cleared: Never

Test-profile Configuration:
  Test-profile name: prof_b2b
  Test packet size: 64 1280
  Therotical max bandwidth : 993000 kbps

Test Configuration:
  Test mode: Initiate-and-Terminate
  Duration in seconds: 20
  Test family: INET
  Routing Instance Name: default
  
```

```
Inet family Configuration:
  Egress Interface : ge-0/1/1.0
  Source ipv4 address: 20.6.0.1
  Destination ipv4 address: 20.6.0.2
  Source udp port: 2020
  Destination udp port: 3030
```

```
Rfc2544 Back-Back test information :
  Initial burst length: 20 seconds at 993000 kbps
  Test iteration mode : Binary
  Test iteration step percent : 50.00 %
```

```
Test packet size: 64
Iteration  Burst Length  Elapsed      Tx      Rx      Tx
          Rx
          (Packets)    time      Packets  Packets  Bytes
          Bytes
1          404754        3        404754   404754   27523272
27523272
2          26983680      20       26983680  26983680  1834890240
1834890240
```

```
Result of the iteration runs : Back-Back-Frames Test complete for packet size
64
```

```
Best iteration : 2
Measured burst (num sec) : 20 sec,
Measured burst (num pkts) : 26983680 packets
Result of the iteration runs : Back-Back-Frames Test complete for packet size
64
```

```
Best iteration : 2
Measured burst (num sec) : 20 sec,
Measured burst (num pkts) : 26983680 packets
```

```
Test packet size: 1280
Iteration  Burst Length  Elapsed      Tx      Rx      Tx
          Rx
          (Packets)    time      Packets  Packets  Bytes
          Bytes
1          28467         3        28467    28467    36551628
36551628
2          1897920      20       1897920  1897920  2436929280
2436929280
```

```
Result of the iteration runs : Back-Back-Frames Test complete for packet size
12
```

```
Best iteration : 2
Measured burst (num sec) : 20 sec,
Measured burst (num pkts) : 1897920 packets
```

```
RFC2544 Back-Back test results summary:
```

```
-----
Packet      Measure Burst
Size        length (Packets)
64          26983680 packets
1280        1897920 packets
```

[show services rpm rfc2544-benchmarking test-id detail \(Reflection Test on MX104 routers\)](#)

```
user@host> show services rpm rfc2544-benchmarking test-id detail 1
```



```

Test information :
  Test id: 1, Test name: fort_uni_inet_ref, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: RFC2544_TEST_STATE_RUNNING
  Status: Running
  Test start time: 2013-12-09 16:24:52 IST
  Test finish time: TEST_RUNNING
  Counters last cleared: Never

```

```

Test Configuration:
  Test mode: Reflect
  Duration in seconds: 864000
  Test family: INET
  Routing Instance Name: default

```

```

Inet family Configuration:
  Egress Interface : ge-0/3/1.0
  Destination ipv4 address: 21.1.1.2
  Destination udp port: 200

```

| Elapsed
time | Reflected
Packets | Reflected
Bytes |
|-----------------|----------------------|--------------------|
| 176 | 8977917 | 9031784502 |

show services rpm rfc2544-benchmarking test-id brief (Reflection Test on MX104 routers)

```

user@host> show services rpm rfc2544-benchmarking test-id brief 1
Test information :
  Test id: 1, Test name: fort_uni_inet_ref, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: RFC2544_TEST_STATE_RUNNING
  Status: Running
  Test start time: 2013-12-09 16:24:52 IST
  Test finish time: TEST_RUNNING
  Counters last cleared: Never

```

show services rpm rfc2544-benchmarking test-id detail (Reflection Test on MX104 routers)

```

user@host> show services rpm rfc2544-benchmarking test-id detail 2
Test information :
  Test id: 2, Test name: fort_uni_inet_ref, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: RFC2544_TEST_STATE_RUNNING
  Status: Running
  Test start time: 2013-12-09 16:39:18 IST
  Test finish time: TEST_RUNNING
  Counters last cleared: Never

Test Configuration:
  Test mode: Reflect
  Duration in seconds: 864000
  Test family: CCC
  Routing Instance Name: default

CCC family Configuration:
  Interface : ge-0/3/2.0
  Test direction: Egress

```

| Elapsed
time | Reflected
Packets | Reflected
Bytes |
|-----------------|----------------------|--------------------|
| 23 | 809137 | 825319740 |

show services rpm rfc2544-benchmarking test-id brief (Reflection Test on MX104 routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 2 brief
Test information :
  Test id: 2, Test name: fort_uni_inet_ref, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: RFC2544_TEST_STATE_RUNNING
  Status: Running
  Test start time: 2013-12-09 16:39:18 IST
  Test finish time: TEST_RUNNING
  Counters last cleared: Never
```

show services rpm twamp server connection

| | |
|---------------------------------|--|
| Syntax | show services rpm twamp server connection
<i><connection-id></i> |
| Release Information | Command introduced in Junos OS Release 9.3. |
| Description | Display information about the connections established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control-clients. By default, all established sessions are displayed, unless you specify a session ID when you issue the command. |
| Options | <i>connection-id</i> —(Optional) Display only information about the specified connection ID. |
| Required Privilege Level | view |
| List of Sample Output | show services rpm twamp server connection on page 2281 |
| Output Fields | Table 151 on page 2281 lists the output fields for the show services rpm twamp server connection command. Output fields are listed in the approximate order in which they appear. |

Table 151: show services rpm twamp server connection Output Fields

| Field Name | Field Description |
|----------------|---|
| Connection ID | Connection ID that uniquely identifies the connection between the TWAMP server and a particular client. |
| Client address | Client IP address. |
| Client port | Client port number. |
| Server address | Server IP address. |
| Server port | Server port number. |
| Session count | Session count. |
| Auth mode | Authentication mode. |

Sample Output

show services rpm twamp server connection

```

user@host> show services rpm twamp server connection
  Connection  Client      Client  Server      Server  Session  Auth
   ID         address      port    address     port    count    mode
         4  1.1.1.1      12345  192.168.219.203    890      16    none

```

| | | | | | | | |
|---------------|-----|-----------------|-------|-----------|-------|----|------|
| | 78 | 3.22.1.55 | 345 | 22.2.2.2 | 89022 | 5 | none |
| | 234 | 192.168.219.203 | 2345 | 2.2.22.2 | 3333 | 16 | none |
| | 5 | 221.4.1.1 | 82345 | 2.2.2.2 | 45909 | 16 | |
| authenticated | 1 | 192.168.1.1 | 645 | 32.2.2.23 | 2394 | 16 | |
| encrypted | | | | | | | |

show services rpm twamp server session

| | |
|---------------------------------|---|
| Syntax | show services rpm twamp server session
<i><session-id></i> |
| Release Information | Command introduced in Junos OS Release 9.3. |
| Description | Display information about the sessions established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default, all established sessions are displayed, unless you specify a session ID when you issue the command. |
| Options | <i>session-id</i> —(Optional) Display only information about the specified session ID. |
| Required Privilege Level | view |
| List of Sample Output | show services rpm twamp server session on page 2283 |
| Output Fields | Table 152 on page 2283 lists the output fields for the show services rpm twamp server session command. Output fields are listed in the approximate order in which they appear. |

Table 152: show services rpm twamp server session Output Fields

| Field Name | Field Description |
|--------------------------|---|
| Session ID | Session ID that uniquely identifies the session between the TWAMP server and a particular client. |
| Connection ID | Connection ID that uniquely identifies the connection between the TWAMP server and a particular client. |
| Sender address | Sender IP address. |
| Sender port | Sender port number. |
| Reflector address | Reflector IP address. |
| Reflector port | Reflector port number. |

Sample Output

show services rpm twamp server session

```

user@host> show services rpm twamp server session
  Session  Connection  Sender      Sender  Reflector  Reflector
   ID      ID          address    port    address    port
  ----
      4         44    1.1.1.1    12345   192.168.219.203    890
      78         44    3.22.1.55    345    22.2.2.2    89022
     234        423   192.168.219.203    2345    2.2.22.2    3333
      5         423   221.4.1.1    82345    2.2.2.2    45909
      1         423   192.168.1.1    645    32.2.2.23    2394

```


show services video-monitoring mdi errors fpc-slot

| | |
|---------------------------------|---|
| Syntax | show services video-monitoring mdi errors fpc-slot <i>fpc-slot</i> |
| Release Information | Command introduced in Junos OS Release 14.1. |
| Description | Display video monitoring error statistics. |
| Options | <i>fpc-slot</i> —Number of the fpc slot for which statistics are displayed. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • Inline Video Monitoring Overview on page 1043 |
| List of Sample Output | show services video-monitoring mdi errors fpc-slot on page 2285 |
| Output Fields | Table 153 on page 2285 lists the output fields for the show services video-monitoring mdi errors fpc-slot <i>fpc-slot</i> command. Output fields are listed in the approximate order in which they appear. |

Table 153: show services video-monitoring mdi errors fpc-slot Output Fields

| Field Name | Field Description |
|---------------------------------|--|
| FPC slot | Slot number of the monitored FPC. |
| Flow Insert Error | Number of errors during new flow insert operations. |
| Flow Policer Drops | <p>Number of packets dropped by flow policer process.</p> <p>NOTE: New flows usually arrive within a very short time interval (1.5 microseconds). These errors do not represent the loss of entire flows, because subsequent packets in the flow can establish the flow. All packets are monitored after a flow has been established. Packet forwarding occurs independently of the video monitoring, and packets are not dropped due to video monitoring errors.</p> |
| Unsupported Media Packets Count | Number of packets dropped because they are not media packets or they are unsupported media packets. |
| PID Limit Exceeded | <p>Number of packets unmonitored because the process identifier (PID) limit exceeded has been exceeded.</p> <p>NOTE: The current PID limit is 6.</p> |

Sample Output

show services video-monitoring mdi errors fpc-slot

```
user@host> show services video-monitoring mdi errors fpc-slot 2
```

MDI Errors Information

FPC Slot: 2

Flow Insert Error: 0, Flow Policer Drops: 0

Unsupported Media Packets Count: 0, PID Limit Exceeded: 202995

show services video-monitoring mdi flows fpc-slot

| | |
|---------------------------------|---|
| Syntax | <pre>show services video-monitoring mdi flows fpc-slot <i>fpc-slot</i> <brief> <count> <destination-address> <destination-port> <detail> <input> <interface-name> <output> <rtp> <source-address> <source-port> <template-name> <udp></pre> |
| Release Information | Command introduced in Junos OS Release 14.1. |
| Description | Display inline video monitoring flow statistics. |
| Options | <p>fpc-slot—Number of the slot for which flows are reported.</p> <p>brief—(Optional) Display brief output(default).</p> <p>count—(Optional) Display the number of flows.</p> <p>destination-address—(Optional) Filter output by destination address.</p> <p>destination-port—(Optional) Filter output by destination port.</p> <p>detail—(Optional) Display output in detailed format including media delivery index records.</p> <p>input—(Optional) Filter output by flow direction input.</p> <p>interface-name—(Optional) Filter output by logical interface name.</p> <p>output—(Optional) Filter output by flow direction output.</p> <p>rtp—(Optional) Filter output by flow type rtp.</p> <p>source-address—(Optional) Filter output by source IP address.</p> <p>source-port—(Optional) Filter output by source port.</p> <p>template-name—(Optional) Filter output by media delivery index template name.</p> <p>udp—(Optional) Filter output by flow type MPEG-TS.</p> |
| Required Privilege Level | view |

Related Documentation

- [Inline Video Monitoring Overview on page 1043](#)

List of Sample Output

[show servicesvideo-monitoring mdi flows fpc-slot brief on page 2288](#)
[show services inline-video-monitoring mdi flows detail on page 2289](#)

Output Fields

[Table 154 on page 2288](#) lists the output fields for the **show services inline-video-monitoring mdi flows fpc-slot fpc-slot** command. Output fields are listed in the approximate order in which they appear.

Table 154: show services mdi flows Output Fields

| Field Name | Field Description |
|---------------|--|
| SIP | Source IP address |
| DIP | Destination IP address |
| SP | Source port |
| DP | Destination port |
| Di | Direction (I=Input, O=Output) |
| Ty | Type of flow |
| Last DF:MLR | Delay factor and media loss rate value of last media delivery index record |
| Avg DF:MLR | Average value of delay factor and media loss rate |
| Last MRV | Media rate variation value of last media delivery index record |
| Avg MRV | Average value of media rate variation |
| IFL | Interface name on which flow is receiving |
| Template Name | Name of template associated with flow |

Sample Output

[show servicesvideo-monitoring mdi flows fpc-slot brief](#)

```
user@host> show services inline-video-monitoring mdi flows fpc-slot 2 brief
```

| Sno | SIP | SP | DIP | DP | Di | Ty | Last DF:MLR | Avg |
|------------|----------|----------|----------|------|------------|-----|---------------|-----|
| DF:MLR | | Last MRV | Avg MRV | IFL | | | Template Name | |
| 1 | 20.0.0.2 | 1024 | 30.0.0.2 | 2048 | I | UDP | 70.90:1 | |
| 92.15:8205 | | -7.09 | -9.36 | | xe-2/2/1.0 | | t1 | |

Sample Output

show services inline-video-monitoring mdi flows detail

```
user@host> show services inline-video-monitoring flows fpc-slot 2 detail count 19
```

Format for RTP flows:

```
Source Address: 20.0.0.2, Source Port: 1024
Destination Address: 30.0.0.2, Destination Port: 2048
Last DF:MLR: 3.58:0, Avg DF:MLR: 3.60:0
Last MRV: 0.00, Avg MRV: 0.00
Interface Name: xe-2/2/1.0, Template Name: t1
Flow Direction: Input, Flow Type: RTP, MDI Records Count: 10
```

| Rec No | DF | MLR | MRV |
|--------|------|-----|------|
| 1 | 3.58 | 0 | 0.00 |
| 2 | 3.62 | 0 | 0.00 |
| 3 | 3.59 | 0 | 0.00 |
| 4 | 3.63 | 0 | 0.00 |
| 5 | 3.60 | 0 | 0.00 |
| 6 | 3.64 | 0 | 0.00 |
| 7 | 3.61 | 0 | 0.00 |
| 8 | 3.57 | 0 | 0.00 |
| 9 | 3.62 | 0 | 0.00 |
| 10 | 3.58 | 0 | 0.00 |

Format for MPEG2-TS over UDP flows:

```
Source Address: 20.0.0.2, Source Port: 1024
Destination Address: 30.0.0.2, Destination Port: 2048
Last DF:MLR: 3.63:0, Avg DF:MLR: 3.61:4097
Last MRV: 0.00, Avg MRV: 0.00
Interface Name: xe-2/2/1.0, Template Name: t1
Flow Direction: Input, Flow Type: UDP, MDI Records Count: 10
```

| Rec No | DF | MLR | MRV | PID-0 | PID-1 | PID-2 |
|--------|--------|-------|------|--------|--------|--------|
| | PID-3 | PID-4 | | PID-5 | | |
| MLR | Val | MLR | Val | MLR | Val | MLR |
| 1 | 3.63 | 0 | 0.00 | 0x1f40 | 0 | 0x1f41 |
| 0 | 0x1f54 | 0 | 0x11 | 0 | 0x1020 | 0 |
| 2 | 3.59 | 0 | 0.00 | 0x1f40 | 0 | 0x1f41 |
| 0 | 0x1f54 | 0 | 0x11 | 0 | 0x1020 | 0 |
| 3 | 3.64 | 0 | 0.00 | 0x1f40 | 0 | 0x1f41 |
| 0 | 0x1f54 | 0 | 0x11 | 0 | 0x1020 | 0 |
| 4 | 3.60 | 0 | 0.00 | 0x1f40 | 0 | 0x1f41 |
| 0 | 0x1f54 | 0 | 0x11 | 0 | 0x1020 | 0 |
| 5 | 3.64 | 0 | 0.00 | 0x1f40 | 0 | 0x1f41 |
| 0 | 0x1f54 | 0 | 0x11 | 0 | 0x1020 | 0 |
| 6 | 3.61 | 0 | 0.00 | 0x1f40 | 0 | 0x1f41 |
| 0 | 0x1f54 | 0 | 0x11 | 0 | 0x1020 | 0 |
| 7 | 3.57 | 0 | 0.00 | 0x1f40 | 0 | 0x1f41 |
| 0 | 0x1f54 | 0 | 0x11 | 0 | 0x1020 | 0 |
| 8 | 3.62 | 0 | 0.00 | 0x1f40 | 0 | 0x1f41 |
| 0 | 0x1f54 | 0 | 0x11 | 0 | 0x1020 | 0 |

| | | | | | | | | |
|----|--------|-------|------|--------|--------|--------|---|------|
| 9 | 3.58 | 40977 | 0.00 | 0x1f40 | 40977 | 0x1f41 | 0 | 0x12 |
| 0 | 0x1f54 | 0 | 0x11 | 0 | 0x1020 | 0 | | |
| 10 | 3.63 | 0 | 0.00 | 0x1f40 | 0 | 0x1f41 | 0 | 0x12 |
| 0 | 0x1f54 | 0 | 0x11 | 0 | 0x1020 | 0 | | |

show services video-monitoring mdi stats fpc-slot

| | |
|---------------------------------|--|
| Syntax | show services video-monitoring mdi stats fpc-slot <i>fpc-slot</i> |
| Release Information | Command introduced in Junos OS Release 14.1. |
| Description | Display inline video monitoring statistics. |
| Options | <i>fpc-slot</i> —Number of the fpc slot for which statistics are displayed. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • Inline Video Monitoring Overview on page 1043 |
| List of Sample Output | show services video-monitoring mdi stats fpc-slot on page 2292 |
| Output Fields | Table 155 on page 2291 lists the output fields for the show services video-monitoring mdi stats fpc-slot <i>fpc-slot</i> command. Output fields are listed in the approximate order in which they appear. |

Table 155: show services video-monitoring mdi stats fpc-slot Output Fields

| Field Name | Field Description |
|----------------------|--|
| FPC Slot | Slot number of the monitored FPC |
| Active Flows | Number of active flows currently monitored.
active flows = inserted flows - deleted flows. |
| Total Inserted Flows | Number of flows initiated under video monitoring. |
| Total Deleted Flows | Number of flows deleted due to inactivity timeout. |
| Total Packets Count | Number of total packets monitored. |
| Total Bytes Count | Number of total bytes monitored. |
| DF Alarm Count | Number of delay factor alarms at each of the following levels: <ul style="list-style-type: none"> • Info level • Warning level • Critical level |

Table 155: show services video-monitoring mdi stats fpc-slot Output Fields (*continued*)

| Field Name | Field Description |
|------------------------|--|
| MLR Alarm Count | Number of media loss rate (MLR) alarms at each of the following levels: <ul style="list-style-type: none">• Info level• Warning level• Critical level |
| MRV alarm count | Number of media rate variation (MRV) alarms at each of the following levels: <ul style="list-style-type: none">• Info level• Warning level• Critical level |

Sample Output

show services video-monitoring mdi stats fpc-slot

```
user@host> show services video-monitoring mdi stats fpc-slot 2
MDI Stats Information
FPC Slot: 2
Active Flows: 1, Total Inserted Flows: 1, Total Deleted Flows: 0
Total Packets Count: 746284, Total Bytes Count: 1013453672
DF alarm count: 0, Info level: 0, Warning level: 0, Critical level: 0
MLR alarm count: 0, Info level: 0, Warning level: 0, Critical level: 0
MRV alarm count: 0, Info level: 0, Warning level: 0, Critical level: 0
```

test services rpm rfc2544-benchmarking test

Syntax test services rpm rfc2544-benchmarking test(ACX Series)
 <clear-counters>
 <routing-instance>
 <test-name>
 <test-id>
 <start>>
 <stop>

Syntax test services rpm rfc2544-benchmarking test(MX104 Router)
 <test-name>
 <test-id>
 <start>>
 <stop>

Release Information Command introduced in Junos OS Release 12.3X52 for ACX Series routers.
 Command introduced in Junos OS Release 13.3R1 for MX104 3D Universal Edge Routers.

Description Start or stop an RFC 2544-based benchmarking test. You can start or stop all of the test names that are defined on a router, or start or stop a specific test name. You can also stop a test based on its test identifier. You can also clear the statistical counters associated with the test. When you trigger an RFC 2544-based benchmarking test, it passes through a series of states. These states are displayed in the Test state field in the brief or displayed output information of the **show services rpm rfc2544-benchmarking** command.



NOTE: The RFC 2544 test is stopped at the initiator automatically after the test successfully completes all of the test steps. You need not explicitly enter the **test services rpm rfc2544-benchmarking test <test-name | test-id> stop** command. However, at the reflector, you must explicitly enter this command to stop the test after the test is completed at the initiator.

Options **start**—Start the RFC 2544-based benchmarking test

stop—Terminate the RFC 2544-based benchmarking test

clear-counters—(ACX Series routers only) Clear the statistics associated with the benchmarking test that was run.

routing-instance—(ACX Series routers only) Name of the routing instance for the test.

test-name—Name of the benchmarking test that must be started or stopped.

test-id—Unique identifier of the test that must be stopped. You can stop a test based on the test identifier. You can use the **test-id** option with only the **test services rpm rfc2544-benchmarking stop** command.

Additional Information The test session is supported in out-of-service mode for the underlying service. You must not transmit any traffic to the UNI port, configured as a generator or a reflector, that is being tested during the duration of the test.

Required Privilege Level view

Related Documentation

- [Configuring an RFC 2544-Based Benchmarking Test on page 989](#)
- [RFC2544-Based Benchmarking Tests Overview on page 983](#)
- [rfc2544-benchmarking on page 1740](#)

List of Sample Output [test services rpm rfc2544-benchmarking on page 2294](#)

Output Fields To display the results of the benchmarking test, use the **show services rpm rfc2544-benchmarking** command.

Sample Output

test services rpm rfc2544-benchmarking

```
user@host> test services rpm rfc2544-benchmarking test test-name test1 start
Test "test1" id 56 started
```

The response specifies that a test has been started with test id 56. The test ID can be further used in **show** commands to view test output.

Tunnel and Encryption Services Operational Commands

- [clear ike security-associations](#)
- [clear ipsec security-associations](#)
- [request ipsec switch](#)
- [request security certificate \(signed\)](#)
- [request security certificate \(unsigned\)](#)
- [request security key-pair](#)
- [request system certificate add](#)
- [show ike security-associations](#)
- [show interfaces \(Encryption\)](#)
- [show interfaces \(GRE\)](#)
- [show interfaces \(IP-over-IP\)](#)
- [show interfaces \(Logical Tunnel\)](#)
- [show interfaces \(Multicast Tunnel\)](#)
- [show interfaces \(PIM\)](#)
- [show interfaces \(Virtual Loopback Tunnel\)](#)
- [show ipsec certificates](#)

- `show ipsec redundancy`
- `show ipsec security-associations`
- `show system certificate`

clear ike security-associations

| | |
|---------------------------------|--|
| Syntax | clear ike security-associations
<destination-ip-address> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (Encryption interface on M Series and T Series routers only) Clear information about the current Internet Key Exchange (IKE) security association. This command is valid for dynamic security associations only. For IKEv2, this command creates new security associations for IKE SA and IPSEC SAs. |
| Options | none —Clear all IKE security associations.

destination-ip-address —(Optional) Clear the IKE security association at the specified destination address. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show ike security-associations on page 2305 |
| List of Sample Output | clear ike security-associations on page 2296 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear ike security-associations

```
user@host> clear ike security-associations
```

clear ipsec security-associations

| | |
|---------------------------------|--|
| Syntax | <code>clear ipsec security-associations</code>
<code><sa-name></code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (Encryption interface on M Series and T Series routers only) Clear information about the current IP Security (IPsec) security association. This command is valid for dynamic security associations only. For IKEv1, this command creates new security associations for IKE SA and IPSEC SAs. |
| Options | none —Clear all IPsec security associations.

sa-name —(Optional) Clear the specified security association. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show ipsec security-associations on page 2350 |
| List of Sample Output | clear ipsec security-associations on page 2297 |
| Output Fields | See the show ipsec security-associations for an explanation of output fields. |

Sample Output

clear ipsec security-associations

The following output from the **show ipsec security-associations detail** command is displayed before and after the **clear ipsec security-associations** command is issued:

```

user@host> show ipsec security-associations detail
Security association: sa-dynamic, Interface family: Up

Direction: inbound, SPI: 242379418, State: Installed
Mode: tunnel, Type: dynamic
Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
Soft lifetime: Expires in 22979 seconds
Hard lifetime: Expires in 28739 seconds

Direction: outbound, SPI: 368592771, State: Installed
Mode: tunnel, Type: dynamic
Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
Soft lifetime: Expires in 22979 seconds
Hard lifetime: Expires in 28739 seconds

user@host> clear ipsec security-associations

user@host> show ipsec security-associations detail
Security association: sa-dynamic, Interface family: Up

Direction: inbound, SPI: 1031597683, State: Installed

```

Mode: tunnel, Type: dynamic
Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
Soft lifetime: Expires in 23037 seconds
Hard lifetime: Expires in 28797 seconds

Direction: outbound, SPI: 1618419878, State: Installed
Mode: tunnel, Type: dynamic
Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
Soft lifetime: Expires in 23037 seconds
Hard lifetime: Expires in 28797 seconds

request ipsec switch


| | |
|---------------------------------|---|
| Syntax | <code>request ipsec switch (interface <es-fpc/pic/port> security-associations <sa-name>)</code> |
| Release Information | Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | (Encryption interface on M Series, PTX Series, and T Series routers and EX Series switches only) Manually switch from the primary to the backup encryption services interface, or switch from the primary to the backup IP Security (IPsec) tunnel. |
| Options | <code>interface <es-fpc/pic/port></code> —Switch to the backup encryption interface.
<code>security-associations <sa-name></code> —Switch to the backup tunnel. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show ipsec redundancy on page 2348 |
| List of Sample Output | request ipsec switch on page 2299 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request ipsec switch

```
user@host> request ipsec switch security-associations sa-private
```

request security certificate (signed)

| | |
|---|--|
| Syntax | <code>request security certificate enroll filename <i>filename</i> subject <i>subject</i>
alternative-subject <i>alternative-subject</i> certification-authority <i>certification-authority</i> encoding
(binary pem) key-file <i>key-file</i> domain-name <i>domain-name</i></code> |
| Release Information | Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | (Encryption interface on M Series and T Series routers and EX Series switches only)
Obtain a signed certificate from a certificate authority (CA). The signed certificate validates the CA and the owner of the certificate. The results are saved in a specified file to the <code>/var/etc/ikecert</code> directory. |
| <div> NOTE: For FIPS mode, the digital security certificates must be compliant with the National Institute of Standards and Technology (NIST) SP 800-131A standard. The <code>request security key-pair</code> command is deprecated and not available with Junos in FIPS mode because it generates RSA and DSA keys with sizes of 512 and 1024 bits that are not compliant with the NIST SP 800-131A standard.</div> | |
| Options | <p>filename <i>filename</i>—File that stores the certificate.</p> <p>subject <i>subject</i>—Distinguished name (dn), which consists of a set of components—for example, an organization (o), an organization unit (ou), a country (c), and a locality (l).</p> <p>alternative-subject <i>alternative-subject</i>—Tunnel source address.</p> <p>certification-authority <i>certification-authority</i>—Name of the certificate authority profile in the configuration.</p> <p>encoding (binary pem)—File format used for the certificate. The format can be a binary file or privacy-enhanced mail (PEM), an ASCII base64-encoded format. The default format is binary.</p> <p>key-file <i>key-file</i>—File containing a local private key.</p> <p>domain-name <i>domain-name</i>—Fully qualified domain name.</p> |
| Required Privilege Level | maintenance |
| List of Sample Output | request security certificate (signed) on page 2301 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request security certificate (signed)

```
user@host> request security certificate enroll filename host.crt subject c=uk,o=london
alternative-subject 10.50.1.4 certification-authority verisign key-file host-1.prv domain-name
host.juniper.net
CA name: juniper.net CA file: ca_verisign
local pub/private key pair: host.prv
subject: c=uk,o=london domain name: host.juniper.net
alternative subject: 10.50.1.4
Encoding: binary
Certificate enrollment has started. To view the status of your enrollment, check
the key management process (kmd) log file at /var/log/kmd. <-----
```

request security certificate (unsigned)


| | |
|---------------------------------|--|
| Syntax | <code>request security certificate enroll filename <i>filename</i> ca-file <i>ca-file</i> ca-name <i>ca-name</i> encoding (binary perm) url <i>url</i></code> |
| Release Information | Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | (Encryption interface on M Series and T Series routers and EX Series switches only)
Obtain a certificate from a certificate authority (CA). The results are saved in a specified file to the <code>/var/etc/ikecert</code> directory. |
| Options | <p>filename <i>filename</i>—File that stores the public key certificate.</p> <p>ca-file <i>ca-file</i>—Name of the certificate authority profile in the configuration.</p> <p>ca-name <i>ca-name</i>—Name of the certificate authority.</p> <p>encoding (binary pem)—File format used for the certificate. The format can be a binary file or privacy-enhanced mail (PEM), an ASCII base64-encoded format. The default value is binary.</p> <p>url <i>url</i>—Certificate authority URL.</p> |
| Required Privilege Level | maintenance |
| List of Sample Output | request security certificate (unsigned) on page 2302 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request security certificate (unsigned)

```
user@host> request security certificate enroll filename ca_verisign ca-file verisign ca-name
juniper.net urlxyzcompany URL
http://<verisign ca-name xyzcompany url>/cgi-bin/pkiclient.exe CA name: juniper.net
CA file: verisign Encoding: binary
Certificate enrollment has started. To view the status of your enrollment, check
the key management process (kmd) log file at /var/log/kmd. <-----
```


request security key-pair

| | |
|---|---|
| Syntax | <code>request security key-pair <i>filename</i></code>
<code><size <i>key-size</i>></code>
<code><type (rsa dsa)></code> |
| Release Information | Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | (Encryption interface on M Series and T Series routers and EX Series switches only)
Generate a public and private key pair for a digital certificate. |
| <div>  <p>NOTE: The <code>request security-certificates</code> command is deprecated and are not available with Junos in FIPS mode because security certificates are not compliant with the NIST SP 800-131A standard.</p> </div> | |
| Options | <p><i>filename</i>—Name of a file in which to store the key pair.</p> <p><i>size key-size</i>—(Optional) Key size, in bits. The key size can be 512, 1024, or 2048. The default value is 1024.</p> <p><i>type</i>—(Optional) Algorithm used to encrypt the key:</p> <ul style="list-style-type: none"> • rsa—RSA algorithm. This is the default. • dsa—Digital signature algorithm with Secure Hash Algorithm (SHA). |
| Required Privilege Level | maintenance |
| List of Sample Output | request security key-pair on page 2303 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request security key-pair

```
user@host> request security key-pair security-key-file
```

request system certificate add

| | |
|---------------------------------|--|
| Syntax | <code>request system certificate add (<i>filename</i> terminal)</code> |
| Release Information | Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 11.1 for the QFX Series. |
| Description | (Encryption interface on M Series and T Series routers, PTX Series, and QFX Series switches only) Add a certificate provided by the Juniper Networks certificate authority (CA). |
| Options | <i>filename</i> —Filename (URL, local, or remote).
<i>terminal</i> —Use login terminal. |
| Required Privilege Level | maintenance |
| List of Sample Output | request system certificate add on page 2304 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

request system certificate add

```
user@host> request system certificate add terminal
```

show ike security-associations

| | |
|---------------------------------|---|
| Syntax | show ike security-associations
<brief detail>
<peer-address> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (Encryption interface on M Series and T Series routers only) Display information about Internet Key Exchange (IKE) security associations. |
| Options | <p>none—Display standard information about all IKE security associations.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>peer-address—(Optional) Display IKE security associations for the specified peer address.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear ike security-associations on page 2296 |
| List of Sample Output | show ike security-associations on page 2308
show ike security-associations detail on page 2308 |
| Output Fields | Table 156 on page 2305 lists the output fields for the show ike security-associations command. Output fields are listed in the approximate order in which they appear. |

Table 156: show ike security-associations Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------|---|-----------------|
| IKE peer | Remote end of the IKE negotiation. | detail |
| Role | Part played in the IKE session. The router triggering the IKE negotiation is the initiator, and the router accepting the first IKE exchange packets is the responder. | detail |
| Remote Address | Responder's address. | none specified |
| State | State of the IKE security association: <ul style="list-style-type: none"> • Matured—The IKE security association is established. • Not matured—The IKE security association is in the process of negotiation. | none specified |
| Initiator cookie | When the IKE negotiation is triggered, a random number is sent to the remote node. | All levels |

Table 156: show ike security-associations Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------------|---|-----------------|
| Responder cookie | <p>The remote node generates its own random number and sends it back to the initiator as a verification that the packets were received.</p> <p>Of the numerous security services available, protection against denial of service (DoS) is one of the most difficult to address. A “cookie” or anticlogging token (ACT) is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity. An exchange prior to CPU-intensive public key operations can thwart some DoS attempts (such as simple flooding with invalid IP source addresses).</p> | All levels |
| Exchange type | <p>Specifies the number of messages in an IKE exchange, and the payload types that are contained in each message. Each exchange type provides a particular set of security services, such as anonymity of the participants, perfect forward secrecy of the keying material, and authentication of the participants. Junos OS supports two types of exchanges:</p> <ul style="list-style-type: none"> • Main—The exchange is done with six messages. Main encrypts the payload, protecting the identity of the neighbor. • Aggressive—The exchange is done with three messages. Aggressive does not encrypt the payload, leaving the identity of the neighbor unprotected. | All Levels |
| Authentication method | Type of authentication determines which payloads are exchanged and when they are exchanged. The Junos OS supports only pre-shared keys . | detail |
| Local | Prefix and port number of the local end. | detail |
| Remote | Prefix and port number of the remote end. | detail |
| Lifetime | Number of seconds remaining until the IKE security association expires. | detail |
| Algorithms | <p>Header for the IKE algorithms output.</p> <ul style="list-style-type: none"> • Authentication—Type of authentication algorithm used: md5 or sha1. • Encryption—Type of encryption algorithm used: des-cbc, 3des-cbc, or None. • Pseudo random function—Function that generates highly unpredictable random numbers: hmac-md5 or hmac-sha1. | detail |
| Traffic statistics | <p>Number of bytes and packets received and transmitted on the IKE security association.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the IKE security association. • Input packets, Output packets—Number of packets received and transmitted on the IKE security association. | detail |

Table 156: show ike security-associations Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|---|-----------------|
| Flags | Notification to the key management process of the status of the IKE negotiation: <ul style="list-style-type: none"> • caller notification sent—Caller program notified about the completion of the IKE negotiation. • waiting for done—Negotiation is done. The library is waiting for the remote end retransmission timers to expire. • waiting for remove—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation. • waiting for policy manager—Negotiation is waiting for a response from the policy manager. | detail |
| IPsec security associates | Number of IPsec security associations created and deleted with this IKE security association. | detail |
| Phase 2 negotiations in progress | Number of phase 2 IKE negotiations in progress and status information: <ul style="list-style-type: none"> • Negotiation type—Type of phase 2 negotiation. The Junos OS currently supports quick mode. • Message ID—Unique identifier for a phase 2 negotiation. • Local identity—Identity of the local phase 2 negotiation. The format is <i>id-type-name (proto-name:port-number,[O..id-data-len] = iddata-presentation)</i> • Remote identity—Identity of the remote phase 2 negotiation. The format is <i>id-type-name (proto-name:port-number,[O..id-data-len] = iddata-presentation)</i> • Flags—Notification to the key management process of the status of the IKE negotiation: <ul style="list-style-type: none"> • caller notification sent—Caller program notified about the completion of the IKE negotiation. • waiting for done—Negotiation is done. The library is waiting for the remote end retransmission timers to expire. • waiting for remove—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation. • waiting for policy manager—Negotiation is waiting for a response from the policy manager. | detail |

Sample Output

show ike security-associations

```
user@host> show ike security-associations
Remote Address  State      Initiator cookie  Responder cookie  Exchange type
4.4.4.4         Matured          93870456fa000011 723a20713700003e Main
```

show ike security-associations detail

```
user@host> show ike security-associations detail
IKE peer 4.4.4.4
Role: Initiator, State: Matured
Initiator cookie: cf22bd81a7000001, Responder cookie: fe83795c2800002e
Exchange type: Main, Authentication method: Pre-shared-keys
Local: 4.4.4.5:500, Remote: 4.4.4.4:500
Lifetime: Expires in 187 seconds
Algorithms:
Authentication      : md5
Encryption           : 3des-cbc
Pseudo random function: hmac-md5
Traffic statistics:
Input bytes  :          1000
Output bytes :          1280
Input packets:           5
Output packets:          9
Flags: Caller notification sent
IPsec security associations: 2 created, 0 deleted
Phase 2 negotiations in progress: 1

Negotiation type: Quick mode, Role: Initiator, Message ID: 3582889153
Local: 4.4.4.5:500, Remote: 4.4.4.4:500
Local identity: ipv4_subnet(tcp:80,[0..7]=10.1.1.0/24)
Remote identity: ipv4_subnet(tcp:100,[0..7]=10.1.2.0/24)
Flags: Caller notification sent, Waiting for done
```

show interfaces (Encryption)

| | |
|---------------------------------|--|
| Syntax | <pre>show interfaces es-fpc/pic/port:channel <brief detail extensive terse> <descriptions> <media> <snmp-index snmp-index> <statistics></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (M Series and T Series routers only) Display status information about the specified encryption interface. |
| Options | <p>es-fpc/pic/port:channel—Display standard status information about the specified encryption interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index snmp-index—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| List of Sample Output | show interfaces (Encryption) on page 2312
show interfaces brief (Encryption) on page 2312
show interfaces detail (Encryption) on page 2312
show interfaces extensive (Encryption) on page 2313 |
| Output Fields | Table 157 on page 2309 lists the output fields for the show interfaces (ES) command. Output fields are listed in the approximate order in which they appear. |

Table 157: Encryption show interfaces Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|--|------------------------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface's index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |

Table 157: Encryption show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|--|------------------------------|
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Type | Encapsulation being used on the interface. | All levels |
| Link-level type | Encapsulation being used on the physical interface. | All levels |
| MTU | MTU size on the physical interface. | All levels |
| Speed | Speed at which the interface is running. | All levels |
| Device flags | Information about the physical device. Possible values are described in the "Link Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Input rate | Input rate in bits per second (bps) and packets per second (pps). | None specified |
| Output rate | Output rate in bps and pps. | None specified |
| Hold-times | Current interface hold-time up and hold-time down, in milliseconds. | detail extensive |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Traffic statistics | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. • Anti-replay failures—Total number of antireplay failures seen on all tunnels configured on the ES PIC. • Authentication—Total number of authentication failures seen on all tunnels configured on the ES PIC. | detail extensive |
| Egress queues | Total number of egress queues supported on the specified interface. | detail extensive |
| Queue counters | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. | detail extensive |
| Logical Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |

Table 157: Encryption show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------|---|-----------------------|
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Flags | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| IP-Header | IP header of the logical interface. | All levels |
| Encapsulation | Encapsulation on the logical interface. | All levels |
| <i>protocol-family</i> | Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed. | brief |
| Input packets | Number of packets received on the logical interface. | None specified |
| Output packets | Number of packets transmitted on the logical interface. | None specified |
| Traffic statistics | Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Local statistics | Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Transit statistics | Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Protocol | Protocol family configured on the logical interface, such as iso , inet6 , mpls . | detail extensive none |
| MTU | MTU size on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route table | Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 . | detail extensive |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Addresses, Flags | Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . Address | detail extensive none |
| Destination | IP address of the remote side of the connection. | detail extensive none |

Table 157: Encryption show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------|---|-----------------------|
| Local | IP address of the logical interface. | detail extensive none |
| Broadcast | Broadcast address of the logical interface. | detail extensive |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |

Sample Output

show interfaces (Encryption)

```

user@host> show interfaces es-0/3/0
Physical interface: es-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 71
  Type: IPSEC, Link-level type: IPSEC-over-IP, MTU: 3900, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

Logical interface es-0/3/0.0 (Index 70) (SNMP ifIndex 45)
  Flags: Hardware-Down Point-To-Point SNMP-Traps
  IP-Header 10.0.10.2:10.0.10.1::df:64:00000000 Encapsulation: IPSEC
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 3800
  Flags: None
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
  Destination: 10.10.0.2, Local: 10.10.0.1

```

show interfaces brief (Encryption)

```

user@host> show interfaces es-0/3/0 brief
Physical interface: es-0/3/0, Enabled, Physical link is Up
  Type: IPSEC, Link-level type: IPSEC-over-IP, MTU: 3900, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

Logical interface es-0/3/0.0
  Flags: Hardware-Down Point-To-Point SNMP-Traps
  IP-Header 10.0.10.2:10.0.10.1::df:64:00000000 Encapsulation: IPSEC
  inet 10.10.0.1 --> 10.10.0.2s

```

show interfaces detail (Encryption)

```

user@host> show interfaces es-0/3/0 detail
Physical interface: es-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 71, Generation: 21
  Type: IPSEC, Link-level type: IPSEC-over-IP, MTU: 3900, Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags    : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps

```

```

Input packets:                0                0 pps
Output packets:               0                0 pps
Anti-replay failures         : 0
Authentication failures      : 0
Egress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets

0 best-effort          0                0                0
1 expedited-fo         0                0                0
2 assured-forw         0                0                0
3 network-cont         0                0                0

```

Logical interface es-0/3/0.0 (Index 70) (SNMP ifIndex 45) (Generation 9)

Flags: Hardware-Down Point-To-Point SNMP-Traps

IP-Header 10.0.10.2:10.0.10.1::df:64:00000000 Encapsulation: IPSEC

Traffic statistics:

```

Input bytes :                0
Output bytes :                0
Input packets:               0
Output packets:              0

```

Local statistics:

```

Input bytes :                0
Output bytes :                0
Input packets:               0
Output packets:              0

```

Transit statistics:

```

Input bytes :                0                0 bps
Output bytes :                0                0 bps
Input packets:               0                0 pps
Output packets:              0                0 pps

```

Protocol inet, MTU: 3800, Generation: 22, Route table: 0

Flags: None

Addresses, Flags: Dest-route-down Is-Preferred Is-Primary

Destination: 10.10.0.2, Local: 10.10.0.1, Broadcast: Unspecified,
Generation: 26

show interfaces extensive (Encryption)

user@host> show interfaces es-0/3/0 extensive

Physical interface: es-0/3/0, Enabled, Physical link is Up

Interface index: 138, SNMP ifIndex: 71, Generation: 21

Type: IPSEC, Link-level type: IPSEC-over-IP, MTU: 3900, Speed: 800mbps

Hold-times : Up 0 ms, Down 0 ms

Device flags : Present Running

Interface flags: Point-To-Point SNMP-Traps

Statistics last cleared: Never

Traffic statistics:

```

Input bytes :                0                0 bps
Output bytes :                0                0 bps
Input packets:               0                0 pps
Output packets:              0                0 pps

```

Anti-replay failures : 0

Authentication failures : 0

Egress queues: 4 supported, 4 in use

```

Queue counters:      Queued packets  Transmitted packets  Dropped packets

0 best-effort          0                0                0

```

| | | | |
|----------------|---|---|---|
| 1 expedited-fo | 0 | 0 | 0 |
| 2 assured-forw | 0 | 0 | 0 |
| 3 network-cont | 0 | 0 | 0 |

Logical interface es-0/3/0.0 (Index 70) (SNMP ifIndex 45) (Generation 9)

Flags: Hardware-Down Point-To-Point SNMP-Traps

IP-Header 10.0.10.2:10.0.10.1::df:64:00000000 Encapsulation: IPSEC

Traffic statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Local statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Transit statistics:

Input bytes : 0 0 bps

Output bytes : 0 0 bps

Input packets: 0 0 pps

Output packets: 0 0 pps

Protocol inet, MTU: 3800, Generation: 22, Route table: 0

Flags: None

Addresses, Flags: Dest-route-down Is-Preferred Is-Primary

Destination: 10.10.0.2, Local: 10.10.0.1, Broadcast: Unspecified,

Generation: 26

show interfaces (GRE)


| | |
|---------------------------------|---|
| Syntax | <pre>show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display status information about the specified generic routing encapsulation (GRE) interface. |
| Options | <p><i>interface-type</i>—On M Series and T Series routers and EX Series switches, the interface type is <i>gr-fpc/pic/port</i>.</p> <p><i>brief detail extensive terse</i>—(Optional) Display the specified output level of interface information.</p> <p><i>descriptions</i>—(Optional) Display interface description strings.</p> <p><i>media</i>—(Optional) Display media-specific information about network interfaces.</p> <p><i>snmp-index snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><i>statistics</i>—(Optional) Display static interface statistics.</p> |
| | <div>  <p>NOTE: You can configure generic routing encapsulation (GRE) interfaces (<i>gre-x/y/z</i>) only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information about GMPLS, see the <i>MPLS Applications Feature Guide for Routing Devices</i> and the <i>Junos OS, Release 14.2</i>.</p> </div> |
| Required Privilege Level | view |
| List of Sample Output | <p>show interfaces (GRE) on page 2319</p> <p>show interfaces brief (GRE) on page 2319</p> <p>show interfaces detail (GRE) on page 2319</p> <p>show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch on page 2320</p> <p>show interfaces extensive (GRE) on page 2321</p> |
| Output Fields | <p>Table 158 on page 2316 lists the output fields for the show interfaces (GRE) command. Output fields are listed in the approximate order in which they appear.</p> |

Table 158: GRE show interfaces Output Fields

| Field Name | Field Description | Level of Output |
|--------------------------------|--|------------------------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface's index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Type | Type of interface. | All levels |
| Link-level type | Encapsulation used on the physical interface. | All levels |
| MTU | MTU size on the physical interface. | All levels |
| Speed | Speed at which the interface is running. | All levels |
| Hold-times | Current interface hold-time up and hold-time down, in milliseconds. | detail extensive |
| Device Flags | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Interface Flags | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Input rate | Input rate in bits per second (bps) and packets per second (pps). | None specified |
| Output rate | Output rate in bps and pps. | None specified |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Traffic statistics | <p>The number of and the rate at which input and output bytes and packets are received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. | detail extensive |
| Logical Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |

Table 158: GRE show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|--|-----------------------|
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| Flags | <p>Information about the logical interface. Possible values listed in the “Logical Interface Flags” section under <i>Common Output Fields Description</i>. describe general information about the logical interface.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> • Reassemble-Pkts—If the Flags field includes this string, the GRE tunnel is configured to reassemble tunnel packets that were fragmented after tunnel encapsulation. | All levels |
| IP-Header | <p>IP header of the logical interface. If the tunnel key statement is configured, this information is included in the IP Header entry.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> • df—If the IP-Header field includes this string immediately following the 16 bits of identification information (that is, if :df: displays after the twelfth byte), the GRE tunnel is configured to allow fragmentation of GRE packets after encapsulation. | All levels |
| Encapsulation | Encapsulation on the logical interface. | All levels |
| Copy-tos-to-outer-ip-header | <p>Status of type of service (ToS) bits in the GRE packet header:</p> <ul style="list-style-type: none"> • On—ToS bits were copied from the payload packet header into the header of the IP packet sent through the GRE tunnel. • Off—ToS bits were not copied from the payload packet header and are set to 0 in the GRE packet header. <p>NOTE: EX Series switches do not support copying ToS bits to the encapsulated packet, so the value of this field is always Off in switch output.</p> | detail extensive |
| Gre keepalives configured | <p>Indicates whether a GRE keepalive time and hold time are configured for the GRE tunnel.</p> <p>NOTE: EX Series switches do not support configuration of GRE tunnel keepalive times and hold times, so the value of this field is always Off in switch output.</p> | detail extensive |
| Gre keepalives adjacency state | Status of the other end of the GRE tunnel: Up or Down . If keepalive messages are not received by either end of the GRE tunnel within the hold-time period, the GRE keepalive adjacency state is down even when the GRE tunnel is up. | detail extensive |
| Input packets | Number of packets received on the logical interface. | None specified |
| Output packets | Number of packets transmitted on the logical interface. | None specified |

Table 158: GRE show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------|---|------------------------------|
| Traffic statistics | <p>Rate of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input rate—Rate of bits and packets received on the interface. • Output rate—Rate of bits and packets transmitted on the interface. | detail extensive |
| Local statistics | Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Transit statistics | Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive none |
| Protocol | Protocol family configured on the logical interface, such as iso , inet6 , or mpls . | detail extensive none |
| <i>protocol-family</i> | Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed. | brief |
| MTU | MTU size on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route table | Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 . | detail extensive |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Addresses, Flags | Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Destination | IP address of the remote side of the connection. | detail extensive none |
| Local | IP address of the logical interface. | detail extensive none |
| Broadcast | Broadcast address of the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |

Sample Output

show interfaces (GRE)

```

user@host> show interfaces gr-1/2/0
Physical interface: gr-0/0/0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 26
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1476
  Flags: None
  Addresses, Flags: Is-Primary
    Local: 1.10.1.1

```

show interfaces brief (GRE)

```

user@host> show interfaces gr-1/2/0 brief
Physical interface: gr-1/2/0, Enabled, Physical link is Up
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

Logical interface gr-1/2/0.0
  Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000
  IP-Header 10.10.0.2:10.10.0.1:47:df:64:0000000000000000
  Encapsulation: GRE-NULL
  inet 10.100.0.1/30
  mpls

```

show interfaces detail (GRE)

```

user@host> show interfaces gr-1/2/0 detail
Physical interface: gr-0/0/0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 26, Generation: 13
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags    : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   : 0 0 bps
    Output bytes  : 0 0 bps
    Input packets : 0 0 pps
    Output packets: 0 0 pps

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47) (Generation 8)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Traffic statistics:
    Input bytes   : 0
    Output bytes  : 0
    Input packets : 0

```

```

Output packets:                0
Local statistics:
Input bytes :                  0
Output bytes :                  0
Input packets:                 0
Output packets:                0
Transit statistics:
Input bytes :                  0          0 bps
Output bytes :                  0          0 bps
Input packets:                 0          0 pps
Output packets:                0          0 pps
Protocol inet, MTU: 1476, Generation: 12, Route table: 0
Flags: None
Addresses, Flags: Is-Primary
Destination: Unspecified, Local: 1.10.1.1, Broadcast: Unspecified,
Generation: 15

```

show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch

```

user@switch> show interfaces gr-2/0/15 detail
Physical interface: gr-2/0/15, Enabled, Physical link is Up
Interface index: 195, SNMP ifIndex: 846, Generation: 198
Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:1f:12:38:0f:d2, Hardware address: 00:1f:12:38:0f:d2
Device flags    : Present Running
Interface flags: Point-To-Point SNMP-Traps
Statistics last cleared: 2011-09-14 17:43:15 UTC (00:00:18 ago)
Traffic statistics:
Input bytes :          5600636          0 bps
Output bytes :          5600636          0 bps
Input packets:          20007          0 pps
Output packets:          20007          0 pps
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0

Logical interface gr-2/0/15.0 (Index 75) (SNMP ifIndex 847) (HW Token 4093)
(Generation 140)
Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 180.20.30.2:180.20.3:47:df:64:0000000000000000
Encapsulation: GRE-NULL
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Traffic statistics:
Input bytes :          5600886
Output bytes :          2881784
Input packets:          20010
Output packets:          10018
Local statistics:
Input bytes :           398
Output bytes :           264
Input packets:           5
Output packets:           3
Transit statistics:
Input bytes :          5600488          0 bps
Output bytes :          2881520          0 bps
Input packets:          20005          0 pps
Output packets:          10015          0 pps

```

```

Protocol inet, Generation: 159, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 90.90.90/24, Local: 90.90.90.10, Broadcast: 90.90.90.255,
  Generation: 144

```

```

Logical interface gr-2/0/15.1 (Index 80) (SNMP ifIndex 848) (HW Token 4088)
(Generation 150)

```

```

Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 160.20.40.2:160.20.30.1:47:df:64:0000000000000000
Encapsulation: GRE-NULL
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down

```

```
Traffic statistics:
```

```

Input bytes :          260
Output bytes :        2880148
Input packets:           4
Output packets:       10002

```

```
Local statistics:
```

```

Input bytes :          112
Output bytes :           0
Input packets:           2
Output packets:           0

```

```
Transit statistics:
```

```

Input bytes :          148          0 bps
Output bytes :       2880148          0 bps
Input packets:           2          0 pps
Output packets:       10002          0 pps

```

```
Protocol inet, Generation: 171, Route table: 0
```

```
Flags: None
```

```
Addresses, Flags: Is-Preferred Is-Primary
```

```

  Destination: 70.70.70/24, Local: 70.70.70.10, Broadcast: 70.70.70.255,
  Generation: 160

```

show interfaces extensive (GRE)

The output for the **show interfaces extensive** command is identical to that for the **show interfaces detail** command. For sample output, see [show interfaces detail \(GRE\) on page 2319](#) and [show interfaces detail \(GRE\) on an EX4200 Virtual Chassis Member Switch on page 2320](#).

show interfaces (IP-over-IP)

| | |
|---------------------------------|--|
| Syntax | <pre>show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display status information about the specified IP-over-IP interface. |
| Options | <p><i>interface-type</i>—On M Series and T Series routers, the interface type is <i>ip-fpc/pic/port</i>.</p> <p><i>brief detail extensive terse</i>—(Optional) Display the specified level of output.</p> <p><i>descriptions</i>—(Optional) Display interface description strings.</p> <p><i>media</i>—(Optional) Display media-specific information about network interfaces.</p> <p><i>snmp-index snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><i>statistics</i>—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show interfaces (IP-over-IP) on page 2324</p> <p>show interfaces brief (IP-over-IP) on page 2325</p> <p>show interfaces detail (IP-over-IP) on page 2325</p> <p>show interfaces extensive (IP-over-IP) on page 2325</p> |
| Output Fields | <p>Table 159 on page 2322 lists the output fields for the show interfaces (IP-over-IP) command. Output fields are listed in the approximate order in which they appear.</p> |

Table 159: IP-over-IP show interfaces Output Fields

| Field | Field Description | Level of Output |
|---------------------------|--|------------------------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface's index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |

Table 159: IP-over-IP show interfaces Output Fields (*continued*)

| Field | Field Description | Level of Output |
|--------------------------------|---|------------------------------|
| Type | Type of interface. | All levels |
| Link-level type | Encapsulation used on the physical interface. | All levels |
| MTU | MTU size on the physical interface. | All levels |
| Speed | Speed at which the interface is running. | All levels |
| Hold-times | Current interface hold-time up and hold-time down, in milliseconds. | detail extensive |
| Device flags | Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Input rate | Input rate in bits per second (bps) and packets per second (pps). | None specified |
| Output rate | Output rate in bps and pps. | None specified |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Traffic statistics | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. | detail extensive |
| Logical Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| Flags | Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> . | All levels |
| IP Header | IP header of the logical interface. | All levels |
| Encapsulation | Encapsulation on the logical interface. | All levels |
| Input packets | Number of packets received on the logical interface. | None specified |

Table 159: IP-over-IP show interfaces Output Fields (*continued*)

| Field | Field Description | Level of Output |
|-------------------------------|---|------------------------------|
| Output packets | Number of packets transmitted on the logical interface. | None specified |
| Traffic statistics | <p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input rate—Rate of bits and packets received on the interface. • Output rate—Rate of bits and packets transmitted on the interface. | detail extensive |
| Local statistics | Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Transit statistics | Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Protocol | Protocol family configured on the logical interface, such as iso , inet6 , or mpls . | detail extensive none |
| <i>protocol-family</i> | Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed. | brief |
| MTU | MTU size on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route table | Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 . | detail extensive |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |

Sample Output

show interfaces (IP-over-IP)

```

user@host> show interfaces ip-0/0/0
Physical interface: ip-0/0/0, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 27
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ip-0/0/0.0 (Index 69) (SNMP ifIndex 49)
    Flags: Point-To-Point SNMP-Traps 16384
    IP-Header 2.2.2.1:2.2.2.2:4:df:64:00000000 Encapsulation: IPv4-NUL
    Input packets : 0

```

```

Output packets: 0
Protocol inet, MTU: 1480
Flags: None

```

show interfaces brief (IP-over-IP)

```

user@host> show interfaces ip-0/0/0 brief
Physical interface: ip-0/0/0, Enabled, Physical link is Up
Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
Device flags : Present Running
Interface flags: SNMP-Traps

Logical interface ip-0/0/0.0
Flags: Point-To-Point SNMP-Traps 16384
IP-Header 2.2.2.1:2.2.2.2:4:df:64:00000000 Encapsulation: IPv4-NULl
inet

```

show interfaces detail (IP-over-IP)

```

user@host> show interfaces ip-0/0/0 detail
Physical interface: ip-0/0/0, Enabled, Physical link is Up
Interface index: 133, SNMP ifIndex: 27, Generation: 14
Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
Hold-times : Up 0 ms, Down 0 ms
Device flags : Present Running
Interface flags: SNMP-Traps
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps

Logical interface ip-0/0/0.0 (Index 69) (SNMP ifIndex 49) (Generation 9)
Flags: Point-To-Point SNMP-Traps 16384
IP-Header 2.2.2.1:2.2.2.2:4:df:64:00000000 Encapsulation: IPv4-NULl
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: 1480, Generation: 13, Route table: 0
Flags: None

```

show interfaces extensive (IP-over-IP)

The output for the show interfaces extensive command is identical to that for the show interfaces detail command. For sample output, see [show interfaces detail \(IP-over-IP\) on page 2325](#).

show interfaces (Logical Tunnel)

| | |
|---------------------------------|--|
| Syntax | <pre>show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display status information about the specified logical tunnel interface. |
| Options | <p><i>interface-type</i>—On M Series and T Series routers, the interface type is <i>lt-fpc/pic/port</i>.</p> <p><i>brief detail extensive terse</i>—(Optional) Display the specified level of output.</p> <p><i>descriptions</i>—(Optional) Display interface description strings.</p> <p><i>media</i>—(Optional) Display media-specific information about network interfaces.</p> <p><i>snmp-index snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><i>statistics</i>—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| List of Sample Output | show interfaces extensive (Logical Tunnel) on page 2330 |
| Output Fields | Table 160 on page 2326 lists the output fields for the show interfaces (logical tunnel) command. Output fields are listed in the approximate order in which they appear. |

Table 160: Logical Tunnel show interfaces Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|--|------------------------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |

Table 160: Logical Tunnel show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|------------------------------|
| Type | Type of interface. Software-Pseudo indicates a standard software interface with no associated hardware device. | All levels |
| Link-level type | Encapsulation used on the physical interface. | All levels |
| MTU | MTU size on the physical interface. | All levels |
| Clocking | Reference clock source: Internal or External when configured. Otherwise, Unspecified . | All levels |
| Speed | Speed at which the interface is running. | All levels |
| Device flags | Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Link type | Type of link. | All levels |
| Link flags | Information about the link. Possible values are described in the "Link Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Physical info | Information about the physical interface. | All levels |
| Hold-times | Current interface hold-time up and hold-time down, in milliseconds. | detail extensive |
| Current address | Configured MAC address. | detail extensive none |
| Hardware address | Hardware MAC address. | detail extensive none |
| Alternate link address | Backup link address. | detail extensive none |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) . | detail extensive none |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Traffic statistics | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> Input bytes, Output bytes—Number of bytes received and transmitted on the interface. Input packets, Output packets—Number of packets received and transmitted on the interface. | detail extensive |

Table 160: Logical Tunnel show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------|---|------------------------------|
| Input errors | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • Resource errors—Sum of transmit drops. | extensive |
| Output errors | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • MTU errors—Number of packets larger than the MTU threshold. • Resource errors—Sum of transmit drops. | extensive |
| Logical Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Flags | Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Encapsulation | Encapsulation on the logical interface. | All levels |

Table 160: Logical Tunnel show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|---|------------------------------|
| Traffic statistics | <p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input bytes—Rate of bytes received on the interface. • Output bytes—Rate of bytes transmitted on the interface. • Input packets—Rate of packets received on the interface. • Output packets—Rate of packets transmitted on the interface. | detail extensive |
| Local statistics | Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Transit statistics | Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Protocol | Protocol family configured on the logical interface, such as iso , inet6 , mpls . | detail extensive none |
| MTU | MTU size on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route table | Route table in which this address exists. For example, Route table:0 refers to inet.0 . | detail extensive |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Addresses, Flags | Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Destination | IP address of the remote side of the connection. | detail extensive none |
| Local | IP address of the logical interface. | detail extensive none |
| Broadcast | Broadcast address of the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |

Sample Output

show interfaces extensive (Logical Tunnel)

```
user@host> show interfaces lt-1/0/0 extensive
Physical interface: lt-1/0/0, Enabled, Physical link is Up
  Interface index: 143, SNMP ifIndex: 70, Generation: 26
  Type: Logical-tunnel, Link-level type: Logical-tunnel, MTU: 0,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Unspecified
  Link flags     : None
  Physical info  : 13
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:90:69:a6:48:7e, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2004-03-03 15:53:52 PST (22:08:46 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :                0                0 bps
    Output bytes  :                0                0 bps
    Input packets :                0                0 pps
    Output packets:                0                0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0

Logical interface lt-1/0/0.0 (Index 66) (SNMP ifIndex 467) (Generation 3024)
  Flags: Point-To-Point SNMP-Traps 16384 DLCI 100 Encapsulation: FR-NLPID
  Traffic statistics:
    Input bytes   :                0
    Output bytes  :                0
    Input packets :                0
    Output packets:                0
  Local statistics:
    Input bytes   :                0
    Output bytes  :                0
    Input packets :                0
    Output packets:                0
  Transit statistics:
    Input bytes   :                0                0 bps
    Output bytes  :                0                0 bps
    Input packets :                0                0 pps
    Output packets:                0                0 pps
  Protocol inet, MTU: 4470, Generation: 7034, Route table: 0
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: Unspecified,
    Generation: 2054
```

show interfaces (Multicast Tunnel)

| | |
|---------------------------------|--|
| Syntax | <pre>show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display status information about the specified multicast tunnel interface and its logical encapsulation and de-encapsulation interfaces. |
| Options | <p><i>interface-type</i>—On M Series and T Series routers, the interface type is <i>mt-fpc/pic/port</i>.</p> <p><i>brief detail extensive terse</i>—(Optional) Display the specified level of output.</p> <p><i>descriptions</i>—(Optional) Display interface description strings.</p> <p><i>media</i>—(Optional) Display media-specific information about network interfaces.</p> <p><i>snmp-index snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><i>statistics</i>—(Optional) Display static interface statistics.</p> |
| Additional Information | The multicast tunnel interface has two logical interfaces: encapsulation and de-encapsulation. These interfaces are automatically created by the Junos OS for every multicast-enabled VPN routing and forwarding (VRF) instance. The encapsulation interface carries multicast traffic traveling from the edge interface to the core interface. The de-encapsulation interface carries traffic coming from the core interface to the edge interface. |
| Required Privilege Level | view |

List of Sample Output [show interfaces \(Multicast Tunnel\) on page 2333](#)
[show interfaces brief \(Multicast Tunnel\) on page 2333](#)
[show interfaces detail \(Multicast Tunnel\) on page 2333](#)
[show interfaces extensive \(Multicast Tunnel\) on page 2333](#)
[show interfaces \(Multicast Tunnel Encapsulation\) on page 2335](#)
[show interfaces \(Multicast Tunnel De-Encapsulation\) on page 2335](#)

Output Fields Table 161 on page 2332 lists the output fields for the **show interfaces** (Multicast Tunnel) command. Output fields are listed in the approximate order in which they appear.

Table 161: Multicast Tunnel show interfaces Output Fields

| Field Name | Field Description | Level of Output |
|--------------------------------|--|------------------------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface's index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Type | Type of interface. | All levels |
| Link-level type | Encapsulation used on the physical interface. | All levels |
| MTU | MTU size on the physical interface. | All levels |
| Speed | Speed at which the interface is running. | All levels |
| Hold-times | Current interface hold-time up and hold-time down, in milliseconds. | detail extensive |
| Device flags | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Input Rate | Input rate in bits per second (bps) and packets per second (pps). | None specified |
| Output Rate | Output rate in bps and pps. | None specified |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |

Table 161: Multicast Tunnel show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|--|-----------------|
| Traffic statistics | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. | All levels |

Sample Output

show interfaces (Multicast Tunnel)

```
user@host> show interfaces mt-1/2/0
Physical interface: mt-1/2/0, Enabled, Physical link is Up
Interface index: 145, SNMP ifIndex: 41
Type: Multicast-GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
Device flags   : Present Running
Interface flags: SNMP-Traps
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
```

show interfaces brief (Multicast Tunnel)

```
user@host> show interfaces mt-1/2/0 brief
Physical interface: mt-1/2/0, Enabled, Physical link is Up
Type: Multicast-GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
Device flags   : Present Running
Interface flags: SNMP-Traps
```

show interfaces detail (Multicast Tunnel)

```
user@host> show interfaces mt-1/2/0 detail
Physical interface: mt-1/2/0, Enabled, Physical link is Up
Interface index: 145, SNMP ifIndex: 41, Generation: 28
Type: Multicast-GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
Hold-times      : Up 0 ms, Down 0 ms
Device flags    : Present Running
Interface flags: SNMP-Traps
Statistics last cleared: Never
Traffic statistics:
  Input bytes   :          170664562          560000 bps
  Output bytes  :          112345376          368176 bps
  Input packets :           2439107           1000 pps
  Output packets:           2439120           1000 pps
```

show interfaces extensive (Multicast Tunnel)

```
user@host> show interfaces mt-1/2/0 extensive
Physical interface: mt-1/2/0, Enabled, Physical link is Up
Interface index: 141, SNMP ifIndex: 529, Generation: 144
Type: Multicast-GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
Hold-times      : Up 0 ms, Down 0 ms
Device flags    : Present Running
Interface flags: SNMP-Traps
Statistics last cleared: Never
```

```
Traffic statistics:
Input bytes :          170664562          560000 bps
Output bytes :         112345376          368176 bps
Input packets:          2439107           1000 pps
Output packets:         2439120           1000 pps
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:             0
```

Logical interface mt-1/2/0.32768 (Index 83) (SNMP ifIndex 556) (Generation 148)

Flags: Point-To-Point SNMP-Traps 0x4000 IP-Header
232.1.1.1:10.0.0.6:47:df:64:0000000800000000 Encapsulation: GRE=NULL

```
Traffic statistics:
Input bytes :          170418430
Output bytes :         112070294
Input packets:          2434549
Output packets:         2435593
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:             0
Local statistics:
Input bytes :              0
Output bytes :             80442
Input packets:              0
Output packets:            1031
Transit statistics:
Input bytes :          170418430          560000 bps
Output bytes :         111989852          368176 bps
Input packets:          2434549           1000 pps
Output packets:         2434562           1000 pps
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:             0
Protocol inet, MTU: 1572, Generation: 182, Route table: 4
Flags: None
Protocol inet6, MTU: 1572, Generation: 183, Route table: 4
Flags: None
```

Logical interface mt-1/2/0.1081344 (Index 84) (SNMP ifIndex 560) (Generation 149)

```
Flags: Point-To-Point SNMP-Traps 0x6000 Encapsulation: GRE=NULL
Traffic statistics:
Input bytes :          246132
Output bytes :          355524
Input packets:           4558
Output packets:          4558
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:             0
Local statistics:
Input bytes :          246132
Output bytes :              0
```



```

Input packets:          4558
Output packets:         0
Transit statistics:
Input bytes :           0          0 bps
Output bytes :         355524      0 bps
Input packets:         0          0 pps
Output packets:        4558        0 pps
IPv6 transit statistics:
Input bytes :           0
Output bytes :          0
Input packets:         0
Output packets:        0
Protocol inet, MTU: Unlimited, Generation: 184, Route table: 4
Flags: None
Protocol inet6, MTU: Unlimited, Generation: 185, Route table: 4
Flags: None

```

show interfaces (Multicast Tunnel Encapsulation)

```

user@host> show interfaces mt-3/1/0.32768
Logical interface mt-3/1/0.32768 (Index 67) (SNMP ifIndex 0)
Flags: Point-To-Point SNMP-Traps 0x4000
IP-Header 239.1.1.1:10.255.70.15:47:df:64:0000000800000000
Encapsulation: GRE-NULL
Input packets : 0
Output packets: 2
Protocol inet, MTU: Unlimited
Flags: None

```

show interfaces (Multicast Tunnel De-Encapsulation)

```

user@host> show interfaces mt-3/1/0.49152
Logical interface mt-3/1/0.49152 (Index 74) (SNMP ifIndex 0)
Flags: Point-To-Point SNMP-Traps 0x6000 Encapsulation: GRE-NULL
Input packets : 0
Output packets: 2
Protocol inet, MTU: Unlimited
Flags: None

```

show interfaces (PIM)

| | |
|---------------------------------|---|
| Syntax | <pre>show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display status information about the specified Protocol Independent Multicast (PIM) de-encapsulation or PIM encapsulation interface, respectively. |
| Options | <p><i>interface-type</i>—On M Series and T Series routers, the PIM de-encapsulation interface type is pd-fpc/pic/port and the PIM encapsulation interface type is pe-fpc/pic/port.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show interfaces (PIM De-Encapsulation) on page 2337</p> <p>show interfaces brief (PIM De-Encapsulation) on page 2338</p> <p>show interfaces detail (PIM De-Encapsulation) on page 2338</p> <p>show interfaces extensive (PIM Encapsulation) on page 2338</p> <p>show interfaces (PIM Encapsulation) on page 2338</p> <p>show interfaces brief (PIM Encapsulation) on page 2338</p> <p>show interfaces detail (PIM Encapsulation) on page 2339</p> <p>show interfaces extensive (PIM Encapsulation) on page 2339</p> |
| Output Fields | Table 162 on page 2336 lists the output fields for the show interfaces (PIM de-encapsulation or encapsulation) command. Output fields are listed in the approximate order in which they appear. |

Table 162: PIM show interfaces Output Fields

| Field Name | Field Description | Level of Output |
|--------------------|---------------------------------|-----------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |

Table 162: PIM show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|------------------------------|
| Enabled | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface's index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Type | Type of interface. | All levels |
| Link-level type | Encapsulation used on the physical interface. | All levels |
| MTU | MTU size on the physical interface. | All levels |
| Speed | Speed at which the interface is running. | All levels |
| Hold-times | Current interface hold-time up and hold-time down, in milliseconds. | detail extensive |
| Device flags | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Interface flags | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> . | All levels |
| Input Rate | Input rate in bits per second (bps) and packets per second (pps). | None specified |
| Output Rate | Output rate in bps and pps. | None specified |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Traffic statistics | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. | detail extensive |

Sample Output

show interfaces (PIM De-Encapsulation)

```

user@host> show interfaces pd-0/0/0
Physical interface: pd-0/0/0, Enabled, Physical link is Up
  Interface index: 130, SNMP ifIndex: 25
  Type: PIMD, Link-level type: PIM-Decapsulator, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running

```

```
Interface flags: SNMP-Traps
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
```

show interfaces brief (PIM De-Encapsulation)

```
user@host> show interfaces pd-0/0/0 brief
Physical interface: pd-0/0/0, Enabled, Physical link is Up
Type: PIMD, Link-level type: PIM-Decapsulator, MTU: Unlimited, Speed: 800mbps
Device flags   : Present Running
Interface flags: SNMP-Traps
```

show interfaces detail (PIM De-Encapsulation)

```
user@host> show interfaces pd-0/0/0 detail
Physical interface: pd-0/0/0, Enabled, Physical link is Up
Interface index: 130, SNMP ifIndex: 25, Generation: 11
Type: PIMD, Link-level type: PIM-Decapsulator, MTU: Unlimited, Speed: 800mbps
Hold-times      : Up 0 ms, Down 0 ms
Device flags    : Present Running
Interface flags: SNMP-Traps
Statistics last cleared: Never
Traffic statistics:
Input bytes      :                      0                0 bps
Output bytes     :                      0                0 bps
Input packets    :                      0                0 pps
Output packets   :                      0                0 pps
```

show interfaces extensive (PIM Encapsulation)

```
user@host> show interfaces pd-0/0/0 extensive
Physical interface: pd-0/0/0, Enabled, Physical link is Up
Interface index: 130, SNMP ifIndex: 25, Generation: 11
Type: PIMD, Link-level type: PIM-Decapsulator, MTU: Unlimited, Speed: 800mbps
Hold-times      : Up 0 ms, Down 0 ms
Device flags    : Present Running
Interface flags: SNMP-Traps
Statistics last cleared: Never
Traffic statistics:
Input bytes      :                      0                0 bps
Output bytes     :                      0                0 bps
Input packets    :                      0                0 pps
Output packets   :                      0                0 pps
```

show interfaces (PIM Encapsulation)

```
user@host> show interfaces pe-0/0/0
Physical interface: pe-0/0/0, Enabled, Physical link is Up
Interface index: 131, SNMP ifIndex: 26
Type: PIME, Link-level type: PIM-Encapsulator, MTU: Unlimited, Speed: 800mbps
Device flags    : Present Running
Interface flags: SNMP-Traps
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
```

show interfaces brief (PIM Encapsulation)

```
user@host> show interfaces pe-0/0/0 brief
Physical interface: pe-0/0/0, Enabled, Physical link is Up
Type: PIME, Link-level type: PIM-Encapsulator, MTU: Unlimited, Speed: 800mbps
Device flags    : Present Running
Interface flags: SNMP-Traps
```

show interfaces detail (PIM Encapsulation)

```
user@host> show interfaces pe-0/0/0 detail
Physical interface: pe-0/0/0, Enabled, Physical link is Up
  Interface index: 131, SNMP ifIndex: 26, Generation: 12
  Type: PIME, Link-level type: PIM-Encapsulator, MTU: Unlimited, Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags    : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes  :                0                0 bps
    Output bytes :                0                0 bps
    Input packets:                0                0 pps
    Output packets:              0                0 pps
```

show interfaces extensive (PIM Encapsulation)

```
user@host> show interfaces pe-0/0/0 extensive
Physical interface: pe-0/0/0, Enabled, Physical link is Up
  Interface index: 131, SNMP ifIndex: 26, Generation: 12
  Type: PIME, Link-level type: PIM-Encapsulator, MTU: Unlimited, Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags    : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes  :                0                0 bps
    Output bytes :                0                0 bps
    Input packets:                0                0 pps
    Output packets:              0                0 pps
```

show interfaces (Virtual Loopback Tunnel)

| | |
|---------------------------------|--|
| Syntax | <pre>show interfaces vt-fpc/pic/port <brief detail extensive terse> <descriptions> <media> <snmp-index snmp-index> <statistics></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display status information about the specified virtual loopback tunnel interface. |
| Options | <p>vt-fpc/pic/port—Display standard information about the specified virtual loopback tunnel interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index snmp-index—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show interfaces (Virtual Loopback Tunnel) on page 2342</p> <p>show interfaces brief (Virtual Loopback Tunnel) on page 2343</p> <p>show interfaces detail (Virtual Loopback Tunnel) on page 2343</p> <p>show interfaces extensive (Virtual Loopback Tunnel) on page 2343</p> |
| Output Fields | Table 163 on page 2340 lists the output fields for the show interfaces (virtual loopback tunnel) command. Output fields are listed in the approximate order in which they appear. |

Table 163: Virtual Loopback Tunnel show interfaces Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|--|------------------------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels |
| Interface index | Physical interface's index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |

Table 163: Virtual Loopback Tunnel show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|------------------------------|
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Type | Type of interface. | All levels |
| Link-level type | Encapsulation used on the physical interface. | All levels |
| MTU | MTU size on the physical interface. | All levels |
| Speed | Speed at which the interface is running. | All levels |
| Hold-times | Current interface hold-time up and hold-time down, in milliseconds. | detail extensive |
| Device flags | Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Input Rate | Input rate in bits per second (bps) and packets per second (pps). | None specified |
| Output Rate | Output rate in bps and pps. | None specified |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Traffic statistics | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. | detail extensive |
| Logical Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Flags | Information about the logical interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> . | All levels |
| Encapsulation | Encapsulation on the logical interface. | All levels |
| Input packets | Number of packets received on the logical interface. | None specified |
| Output packets | Number of packets transmitted on the logical interface. | None specified |

Table 163: Virtual Loopback Tunnel show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------|---|------------------------------|
| Bandwidth | Bandwidth allotted to the logical interface, in kilobytes per second. | All levels |
| Traffic statistics | <p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. | detail extensive |
| Transit statistics | Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| <i>protocol-family</i> | Protocol family configured on the logical interface. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | brief |
| Protocol | Protocol family configured on the logical interface. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| MTU | Maximum transmission unit size on the logical interface. | detail extensive none |
| Maximum labels | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route Table | Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0. | detail extensive |
| Flags | Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |

Sample Output

show interfaces (Virtual Loopback Tunnel)

```

user@host> show interfaces vt-1/2/0
Physical interface: vt-1/2/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 40
  Type: Loopback, Link-level type: Virtual-loopback-tunnel, MTU: Unlimited,
  Speed: 800mbps
  Device flags      : Present Running
  Input rate       : 0 bps (0 pps)
  Output rate      : 0 bps (0 pps)

Logical interface vt-1/2/0.0 (Index 76) (SNMP ifIndex 57)
  Flags: Point-To-Point 16384 Encapsulation: Virtual-loopback-tunnel

```



```

Input packets : 0
Output packets: 0
  Protocol inet, MTU: Unlimited
    Flags: None
  Protocol mpls, MTU: Unlimited, Maximum labels: 3
    Flags: None

```

show interfaces brief (Virtual Loopback Tunnel)

```

user@host> show interfaces vt-1/2/0 brief
Physical interface: vt-1/2/0, Enabled, Physical link is Up
  Type: Loopback, Link-level type: Virtual-loopback-tunnel, MTU: Unlimited,
  Speed: 800mbps
  Device flags   : Present Running

Logical interface vt-1/2/0.0
  Flags: Point-To-Point 16384 Encapsulation: Virtual-loopback-tunnel
  inet
  mpls

```

show interfaces detail (Virtual Loopback Tunnel)

```

user@host> show interfaces vt-1/2/0 detail
Physical interface: vt-1/2/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 40, Generation: 27
  Type: Loopback, Link-level type: Virtual-loopback-tunnel, MTU: Unlimited,
  Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags    : Present Running
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :                0                0 bps
    Output bytes  :                0                0 bps
    Input packets :                0                0 pps
    Output packets:                0                0 pps

Logical interface vt-1/2/0.0 (Index 76) (SNMP ifIndex 57) (Generation 17)
  Flags: Point-To-Point 16384 Encapsulation: Virtual-loopback-tunnel
  Traffic statistics:
    Input bytes   :                0
    Output bytes  :                0
    Input packets :                0
    Output packets:                0
  Transit statistics:
    Input bytes   :                0                0 bps
    Output bytes  :                0                0 bps
    Input packets :                0                0 pps
    Output packets:                0                0 pps
  Protocol inet, MTU: Unlimited, Generation: 33, Route table: 0
    Flags: None
  Protocol mpls, MTU: Unlimited, Maximum labels: 3, Generation: 34, Route table:
0
    Flags: None

```

show interfaces extensive (Virtual Loopback Tunnel)

```

user@host> show interfaces vt-1/2/0 extensive
Physical interface: vt-1/2/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 40, Generation: 27
  Type: Loopback, Link-level type: Virtual-loopback-tunnel, MTU: Unlimited,
  Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms

```

Device flags : Present Running

Statistics last cleared: Never

Traffic statistics:

| | | |
|-----------------|---|-------|
| Input bytes : | 0 | 0 bps |
| Output bytes : | 0 | 0 bps |
| Input packets: | 0 | 0 pps |
| Output packets: | 0 | 0 pps |

Logical interface vt-1/2/0.0 (Index 76) (SNMP ifIndex 57) (Generation 17)

Flags: Point-To-Point 16384 Encapsulation: Virtual-loopback-tunnel

Traffic statistics:

| | |
|-----------------|---|
| Input bytes : | 0 |
| Output bytes : | 0 |
| Input packets: | 0 |
| Output packets: | 0 |

Transit statistics:

| | | |
|-----------------|---|-------|
| Input bytes : | 0 | 0 bps |
| Output bytes : | 0 | 0 bps |
| Input packets: | 0 | 0 pps |
| Output packets: | 0 | 0 pps |

Protocol inet, MTU: Unlimited, Generation: 33, Route table: 0

Flags: None

Protocol mpls, MTU: Unlimited, Maximum labels: 3, Generation: 34, Route table:

0

Flags: None

show ipsec certificates

| | |
|---------------------------------|---|
| Syntax | show ipsec certificates
<brief detail>
<crl <i>crl-name</i> <i>serial-number</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (Encryption interface on M Series and T Series routers only) Display information about the IPsec certificate database. |
| Options | <p>none—Display standard information about all of the entries in the IPsec certificate database.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>crl <i>crl-name</i> <i>serial-number</i>—(Optional) Display information about the entries on the certificate revocation list (CRL) or for the specified serial number. A CRL is a timestamped list identifying revoked certificates. The CRL is signed by a certificate authority (CA) or CRL issuer and made freely available in a public repository. Each revoked certificate is identified in a CRL by its certificate serial number.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear ipsec security-associations on page 2297 |
| List of Sample Output | show ipsec certificates detail on page 2346 |
| Output Fields | Table 164 on page 2345 lists the output fields for the show ipsec certificates command. Output fields are listed in the approximate order in which they appear. |

Table 164: show ipsec certificates Output Fields

| Field Name | Field Description | Level of Output |
|-----------------|---|-----------------|
| Database | Display information about the IPsec certificate database. <ul style="list-style-type: none"> • Total entries—Number of database entries, including entries that are not trusted or that are in the process of being deleted. • Active entries—Number of database entries, excluding entries that are marked as deleted. • Locked entries—Number of statically configured database entries that cannot expire, such as CA certificates that are root or trusted. | All levels |
| Subject | Distinguished name for the certificate for C, O, CN , as described in RFC 3280, <i>Internet x.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile</i> . | All levels |
| ID | Identification number of the database entry. ID is generated by the internal certificate database. | All levels |

Table 164: show ipsec certificates Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------------|---|-----------------|
| References | Reference number the certificate manager has for the particular entry. | detail |
| Serial | Unique serial number assigned to each certificate by the CA. | All levels |
| Flags | State of the certificate. <ul style="list-style-type: none"> • Trusted—Passed validity checks. • Not trusted—Failed validity checks. • Root—Entry is locked and may have been learned through IKE or a locally configured CA certificate. • Non-root—Entry is not locked. • Crl-issuer—Entity issues CRLs. • Non-crl-issuer—Entity does not issue CRLs. | detail |
| Validity period starts | Start time that the certificate is valid, in the format <i>yyyy mon dd, hh:mm:ss GMT</i> . | detail |
| Validity period ends | End time that the certificate is valid, in the format <i>yyyy mon dd, hh:mm:ss GMT</i> . | detail |
| Alternative name information | Auxiliary identity for the certificate: <i>dns-name</i> , <i>email-address</i> , <i>ip-address</i> , or <i>uri</i> (uniform resource identifier). | detail |
| Issuer | Information about the entity that has signed and issued the CRL as described in RFC 2459, <i>Internet X.509 Public Key Infrastructure Certificate and CRL Profile</i> . | detail |

Sample Output

show ipsec certificates detail

```

user@host> show ipsec certificates detail
Database: Total entries: 3 Active entries: 4 Locked entries: 1
Subject: C=us, O=x
  ID: 5, References: 0, Serial: 22314868
  Flags: Trusted Non-root Crl-issuer
  Validity period starts: 2003 Mar 1st, 01:20:42 GMT
  Validity period ends: 2003 Mar 31st, 01:50:42 GMT
  Alternative name information:
    IP address: 10.20.210.1
  Issuer: C=FI, O=Company-ABC, CN=Company ABC class 2

Subject: C=us, O=x
  ID: 4, References: 0, Serial: 22315496
  Flags: Trusted Non-root Crl-issuer
  Validity period starts: 2003 Mar 1st, 01:21:45 GMT
  Validity period ends: 2003 Mar 31st, 01:51:45 GMT
  Alternative name information:
    IP address: 10.20.210.20
  Issuer: C=FI, O=Company-ABC, CN=Company ABC class 2

Subject: C=FI, O=SSH Company-ABC, CN=Company ABC class 2
  ID: 1, References: 1, Serial: 1538512
  Flags: Trusted Root Non-crl-issuer

```

Validity period starts: 2001 Aug 1st, 07:08:32 GMT
Validity period ends: 2004 Aug 1st, 07:08:32 GMT
Alternative name information:
Email address: certifier-support@ssh.com
Issuer: C=FI, O=Company-ABC, CN=Company ABC class 2

show ipsec redundancy

| | |
|---------------------------------|--|
| Syntax | <code>show ipsec redundancy (interface <es-fpc/pic/port> security association <sa-name>)</code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (Encryption interface on M Series and T Series routers only) Display information about IPsec redundancy. |
| Options | <p>interface <es-fpc/pic/port>—Display information about all encryption interfaces, or optionally, about a particular encryption interface.</p> <p>security association <sa-name>—Display information about all remote tunnels, or optionally, about a particular remote tunnel.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> request ipsec switch on page 2299 |
| List of Sample Output | show ipsec redundancy interface on page 2349
show ipsec redundancy security-associations on page 2349 |
| Output Fields | Table 165 on page 2348 lists the output fields for the show ipsec redundancy command. Output fields are listed in the approximate order in which they appear. |

Table 165: show ipsec redundancy Output Fields

| Field Name | Field Description |
|-----------------------------|---|
| Failure counter | Number of times a PIC switched between primary and backup interfaces, or the number of times the tunnel switched between the primary and remote peers since the software has been activated. |
| Primary interface ' | Name of the interface configured to be the primary interface. |
| Backup interface | Name of the interface configured to be the backup interface. |
| State | State of the primary or backup interface can be Active , Offline , or Standby . Both ES PICs are initialized to Offline . For primary and remote peers, State can be Active or Standby . Both peers are in a state of Standby by default (there is not yet a connection between the two peers). |
| Security association | Name of the security association. |
| Local IP | Local IP address. |
| Primary remote IP | IP address of the configured primary remote peer. |
| Backup remote IP | IP address of the configured backup remote peer. |

Sample Output

show ipsec redundancy interface

```
user@host> show ipsec redundancy interface
Failure counter: 0
Primary interface: es-1/3/0, State: Active
Backup interface : es-1/1/0, State: Standby
```

show ipsec redundancy security-associations

```
user@host> show ipsec redundancy security-associations sa-dynamic
Security association: sa-dynamic, Failure counter: 0
Local IP: 4.4.4.4
Primary remote IP: 4.4.4.5, State: Standby
Backup remote IP : 3.3.3.3, State: Standby
```

show ipsec security-associations

| | |
|---------------------------------|--|
| Syntax | show ipsec security-associations
<brief detail>
<sa-name> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display information about the IPsec security associations applied to the local or transit traffic stream. |
| Options | <p>none—Display standard information about all IPsec security associations.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>sa-name—(Optional) Display the specified IPsec security association.</p> |
| Required Privilege Level | view |
| List of Sample Output | show ipsec security-associations brief (Specific Security Association) on page 2352
show ipsec security-associations detail (Specific Security Association) on page 2352 |
| Output Fields | Table 166 on page 2350 lists the output fields for the show ipsec security-associations command. Output fields are listed in the approximate order in which they appear. |

Table 166: show ipsec security-associations Output Fields

| Field Name | Field Description | Level of Output |
|-----------------------------|--|-----------------|
| Security association | Name of the security association. | All levels |
| Interface family | <p>Status of the interface family of the security association. If the interface family field is absent, it is a transport mode security association. The interface family can have one of three options:</p> <ul style="list-style-type: none"> • Up—The security association is referenced in the interface family and the interface family is up. • Down—The security association is referenced in the interface family and the interface family is down. • No reference—The security association is not referenced in the interface family. | All levels |
| Local gateway | Gateway address of the local system. | All levels |
| Remote gateway | Gateway address of the remote system. | All levels |
| Local identity | Prefix and port number of the local end | All levels |
| Remote identity | Prefix and port number of the remote end. | All levels |
| Direction | Direction of the security association: inbound or outbound . | All levels |
| SPI | Value of the security parameter index. | All levels |

Table 166: show ipsec security-associations Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|--|-----------------|
| AUX-SPI | Value of the auxiliary security parameter index. <ul style="list-style-type: none"> When the value is AH or ESP, AUX-SPI is always 0. When the value is AH+ESP, AUX-SPI is always a positive integer. | All levels |
| State | Status of the security association: <ul style="list-style-type: none"> Installed—The security association is installed in the security association database. (For transport mode security associations, the value of State must always be Installed.) Not installed—The security association is not installed in the security association database. | detail |
| Mode | Mode of the security association: <ul style="list-style-type: none"> transport—Protects single host-to-host protections. tunnel—Protects connections between security gateways. | All levels |
| Type | Type of security association: <ul style="list-style-type: none"> manual—Security parameters require no negotiation. They are static, and are configured by the user. dynamic—Security parameters are negotiated by the IKE protocol. Dynamic security associations are not supported in transport mode. | All levels |
| Protocol | Protocol supported: <ul style="list-style-type: none"> transport mode—Supports Encapsulation Security Protocol (ESP) or Authentication Header (AH). tunnel mode—Supports ESP or AH+ESP. | All levels |
| Authentication | Type of authentication used: hmac-md5-96 , hmac-sha1-96 , or None . | detail |
| Encryption | Type of encryption used: des-cbc , 3des-cbc , aes-128-cbc , aes-192-cbc , aes-256-cbc , or None . | detail |
| Soft lifetime
Hard lifetime | (dynamic output only) Each lifetime of a security association has two display options, hard and soft, one of which must be present for a dynamic security association. The hard lifetime specifies the lifetime of the SA. The soft lifetime , which is derived from the hard lifetime, informs the IPsec key management system that the SA is about to expire. This allows the key management system to negotiate a new SA before the hard lifetime expires. <ul style="list-style-type: none"> Expires in seconds seconds—Number of seconds left until the security association expires. Expires in kilobytes kilobytes—Number of kilobytes left until the security association expires. | detail |
| Anti-replay service | State of the service that prevents packets from being replayed: Enabled or Disabled . | detail |

Table 166: show ipsec security-associations Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------|---|-----------------|
| Replay window size | Configured size, in packets, of the antireplay service window: 32 or 64. The antireplay window size protects the receiver against replay attacks by rejecting old or duplicate packets. If the replay window size is 0, the antireplay service is disabled. | detail |

Sample Output

show ipsec security-associations brief (Specific Security Association)

```

user@host> show ipsec security-associations sa-cosmic brief
Security association: sa-cosmic, Interface family: Up
Local gateway: 21.21.1.1, Remote gateway: 21.21.2.1
Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Direction SPI      AUX-SPI    Mode      Type      Protocol
inbound  2908734119  0          tunnel    dynamic   AH
outbound  3494029335  0          tunnel    dynamic   AH

```

show ipsec security-associations detail (Specific Security Association)

```

user@host> show ipsec security-associations sa-cosmic detail
Security association: sa-cosmic, Interface family: Up

Local gateway: 21.21.1.1, Remote gateway: 21.21.2.1
Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Direction: inbound, SPI: 2908734119, AUX-SPI: 0, State: Installed
Mode: tunnel, Type: dynamic
Protocol: AH, Authentication: hmac-md5-96, Encryption: None
Soft lifetime: Expired
Hard lifetime: Expires in 120 seconds
Anti-replay service: Disabled

Direction: outbound, SPI: 3494029335, AUX-SPI: 0, State: Installed
Mode: tunnel, Type: dynamic
Protocol: AH, Authentication: hmac-md5-96, Encryption: None
Soft lifetime: Expired
Hard lifetime: Expires in 120 seconds
Anti-replay service: Disabled

```

show system certificate

| | |
|---------------------------------|---|
| Syntax | <code>show system certificate</code>
<code><certificate-id></code> |
| Release Information | Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | (Encryption interface on M Series, T Series routers, QFX Series, and OCX Series switches only) Display installed certificates signed by the Juniper Networks certificate authority. |
| Options | none —Display all installed certificates signed by the Juniper Networks certificate authority.

certificate-id —(Optional) Display the details of a particular certificate. |
| Required Privilege Level | maintenance |
| List of Sample Output | show system certificate on page 2354
show system certificate (QFX Series) on page 2354 |
| Output Fields | Table 167 on page 2353 lists the output fields for the show system certificate command. Output fields are listed in the approximate order in which they appear. |

Table 167: show system certificate Output Fields

| Field Name | Field Description |
|---------------------------------|---|
| Certificate identifier | Unique identifier associated with a certificate. The certificate identifier is the common name of the subject. |
| Issuer
Subject | Information about the certificate issuer and the distinguished name (DN) of the issuer, respectively: <ul style="list-style-type: none"> • Organization—Name of the owner's organization. • Organizational unit—Name of the owner's department. • Country—Two-character country code in which the owner's system is located. • State—State in the USA in which the owner is using the certificate. • Locality—City in which the owner's system is located. • Common name—Name of the owner of the certificate. • E-mail address—E-mail address of the owner of the certificate. |
| Validity | When a certificate is valid. |
| Signature algorithm | Encryption algorithm applied to the installed certificate. |
| Public key algorithm | Encryption algorithm applied to the public key. |

Sample Output

show system certificate

```
user@host> show system certificate
Certificate identifier: Dallas-v3
Issuer:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas CA,
E-mail address:ca@juniper.net
Subject:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas-v3,
E-mail address:ca@juniper.net
Validity:
Not before: Mar 13 03:23:25 2004 GMT
Not after: Mar 24 03:23:25 2014 GMT
Signature algorithm: sha1WithRSAEncryption
Public key algorithm: dsaEncryption
```

show system certificate (QFX Series)

```
user@host> show system certificate
Certificate identifier: Dallas-v3
Issuer:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas CA,
E-mail address:ca@juniper.net
Subject:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas-v3,
E-mail address:ca@juniper.net
Validity:
Not before: Mar 13 03:23:25 2004 GMT
Not after: Mar 24 03:23:25 2014 GMT
Signature algorithm: sha1WithRSAEncryption
Public key algorithm: dsaEncryption
```

PART 23

Index

- [Index on page 2357](#)

Index

Symbols

| | |
|--|------|
| #, comments in configuration statements..... | lx |
| (), in syntax descriptions..... | lx |
| 6rd flows | |
| statistics..... | 2035 |
| < >, in syntax descriptions..... | lx |
| [], in configuration statements..... | lx |
| { }, in configuration statements..... | lx |
| (pipe), in syntax descriptions..... | lx |

A

| | |
|--|--------------------|
| AACL | |
| action statements..... | 664 |
| applications..... | 662 |
| best-effort application | |
| identification..... | 658, 671, 693, 699 |
| example configuration..... | 666 |
| logging flows..... | 665 |
| match conditions..... | 662 |
| rules..... | 666 |
| statistics | |
| clearing..... | 2083 |
| aacl-fields statement..... | 1539 |
| aacl-statistics-profile statement..... | 1540 |
| accept | |
| action..... | 871, 1051 |
| accounting statement | |
| flow monitoring..... | 1634 |
| usage guidelines..... | 883, 1063 |
| acknowledge-retries statement..... | 1597 |
| usage guidelines..... | 727 |
| acknowledge-timer statement..... | 1597 |
| usage guidelines..... | 727 |
| action-red-differential-delay statement..... | 1598 |
| usage guidelines..... | 727 |
| active flow monitoring | |
| aggregated flows, displaying..... | 2204 |
| available PICs, displaying..... | 2227 |
| CPU usage, displaying..... | 2230 |
| error statistics, displaying..... | 2209 |
| flow statistics, displaying..... | 2213 |

| | |
|--|----------------|
| flows, detailed information, displaying..... | 2218 |
| memory statistics, displaying..... | 2223 |
| packet size distribution, displaying..... | 2225 |
| adaptive services interfaces..... | 1877 |
| status information, displaying..... | 1877 |
| adaptive-services-pics statement..... | 1300 |
| usage guidelines..... | 844 |
| address statement | |
| APPID | |
| usage guidelines..... | 676 |
| application rule..... | 1541 |
| DFC..... | 1635 |
| usage guidelines..... | 850 |
| encryption..... | 1785 |
| usage guidelines..... | 1251 |
| flow monitoring..... | 1635 |
| usage guidelines..... | 871, 1051 |
| interfaces..... | 1266 |
| usage guidelines..... | 45 |
| link services..... | 1598 |
| usage guidelines..... | 721 |
| NAT..... | 1301 |
| voice services..... | 1301 |
| usage guidelines..... | 622 |
| address-allocation statement..... | 1302 |
| address-range statement | |
| NAT..... | 1302 |
| aggregate-export-interval statement..... | 1636 |
| usage guidelines..... | 883, 1063 |
| aggregated flows, displaying..... | 2204 |
| Aggregated Multiservices interfaces | |
| load balancing..... | 1925 |
| aggregation statement..... | 1303 |
| flow monitoring..... | 1637 |
| usage guidelines..... | 358, 898, 1078 |
| alert (system logging severity level)..... | 21, 47, 645 |
| ALGs | |
| configuring..... | 304 |
| supported on the MS-MIC and MS-MPC..... | 322 |
| ALGsJ | |
| default..... | 141, 300 |
| allow-fragmentation statement..... | 1785 |
| usage guidelines..... | 1217 |
| allow-ip-options statement..... | 1304 |
| usage guidelines..... | 334 |
| allow-multicast statement..... | 1305 |
| usage guidelines..... | 38 |
| allow-overlapping-nat-pools statement..... | 1305 |

| | | | |
|--|--------------------|--|--------------|
| allowed-destinations statement..... | 1638 | NAT..... | 1314 |
| usage guidelines..... | 851 | usage guidelines..... | 70 |
| AMS | | stateful firewall..... | 1315 |
| HA..... | 605, 607 | usage guidelines..... | 332 |
| NAT..... | 605, 616 | usage guidelines..... | 515 |
| analyzer-address statement..... | 1638 | application-system-cache-timeout | |
| usage guidelines..... | 841 | statement..... | 1547 |
| analyzer-id statement..... | 1639 | APPID | |
| usage guidelines..... | 841 | usage guidelines..... | 683 |
| anomaly checklist..... | 328 | applications..... | 357 |
| anti-replay-window-size statement..... | 1306, 1307 | example configuration..... | 321 |
| usage guidelines..... | 428, 433 | applications statement | |
| any (system logging severity level)..... | 21, 47, 644 | AACL..... | 1547 |
| any-any match condition | | usage guidelines..... | 662 |
| Ipsec..... | 424 | APPID | |
| app-mapping-timeout statement..... | 1308 | usage guidelines..... | 680 |
| APPID | | application identification..... | 1548 |
| best-effort application | | application-level gateways..... | 1269, 1317 |
| identification..... | 658, 671, 693, 699 | applications hierarchy..... | 1267, 1315 |
| example configuration..... | 688 | CoS..... | 1267, 1316 |
| application statement..... | 1309, 1542, 1543 | IDS..... | 1268, 1316 |
| APPID | | usage guidelines..... | 357 |
| usage guidelines..... | 674 | NAT..... | 1268, 1317 |
| usage guidelines..... | 303 | usage guidelines..... | 70 |
| application-aware-access-list See aacl | | stateful firewall..... | 1269, 1317 |
| application-aware-access-list-fields | | usage guidelines..... | 332 |
| statement..... | 1544 | usage guidelines..... | 515 |
| application-group statement..... | 1545 | applying service set to interface..... | 31 |
| APPID | | archive-sites statement..... | 1639 |
| usage guidelines..... | 680 | usage guidelines..... | 842 |
| application-group-any statement..... | 1545 | AS PIC | |
| AACL | | multicast traffic..... | 38 |
| usage guidelines..... | 662 | redundancy..... | 41, 648, 826 |
| application-groups statement..... | 1546 | associations, clearing..... | 1825 |
| AACL | | asymmetrical routing support | |
| usage guidelines..... | 662 | APPID..... | 685 |
| APPID | | authentication statement..... | 1318 |
| usage guidelines..... | 680 | usage guidelines..... | 388 |
| application-profile statement..... | 1312 | authentication-algorithm statement | |
| usage guidelines..... | 517 | IKE..... | 1319 |
| application-protocol statement..... | 1310 | usage guidelines..... | 406 |
| usage guidelines..... | 304 | IPsec..... | 1320 |
| application-set statement..... | 1313 | usage guidelines..... | 416 |
| usage guidelines..... | 303 | authentication-method statement..... | 1321 |
| application-sets statement | | usage guidelines..... | 406 |
| CoS..... | 1313 | authentication-mode statement | |
| IDS..... | 1314 | RPM..... | 1640 |
| usage guidelines..... | 357 | | |

| | |
|---------------------------------------|-----------|
| automatic statement..... | 1548 |
| APPID | |
| usage guidelines..... | 686 |
| autonomous-system-type statement..... | 1641 |
| usage guidelines..... | 898, 1078 |
| auxiliary-spi statement..... | 1322 |
| usage guidelines..... | 388 |

B

| | |
|--|--------------------|
| backup AS PIC..... | 41 |
| backup Link Services IQ PIC..... | 548 |
| backup-destination statement..... | 1786 |
| usage guidelines..... | 1260 |
| backup-interface statement..... | 1786 |
| usage guidelines..... | 1259 |
| backup-remote-gateway statement..... | 1322 |
| usage guidelines..... | 427 |
| bandwidth | |
| and delay buffer allocation..... | 532 |
| guaranteed..... | 532, 537 |
| basic-nat-pt option | |
| configuring..... | 143 |
| basic-nat44 option | |
| configuring..... | 93 |
| basic-nat66 option | |
| configuring..... | 99 |
| benchmarking test See RFC2544 benchmarking | |
| test, RPM service | |
| best-effort application | |
| identification..... | 658, 671, 693, 699 |
| bgp statement | |
| RPM..... | 1642 |
| braces, in configuration statements..... | lx |
| brackets | |
| angle, in syntax descriptions..... | lx |
| square, in configuration statements..... | lx |
| bundle statement..... | 1323, 1599 |
| usage guidelines..... | 626, 721 |
| by-destination statement..... | 1323 |
| usage guidelines..... | 358 |
| by-pair statement..... | 1324 |
| usage guidelines..... | 358 |
| by-source statement..... | 1325 |
| usage guidelines..... | 358 |
| bypass-traffic-on-exceeding-flow-limits | |
| statement..... | 1325, 1549 |
| bypass-traffic-on-pic-failure statement..... | 1326 |
| usage guidelines..... | 31 |

C

| | |
|--|------------|
| capture-group statement..... | 1643 |
| usage guidelines..... | 849 |
| certificates | |
| for IKE negotiation, displaying..... | 1958 |
| installed, displaying..... | 2353 |
| key pairs, generating..... | 2303 |
| PKI | |
| CA certificates, clearing..... | 1846 |
| CA certificates, displaying..... | 1930 |
| CA certificates, loading manually..... | 1865 |
| certificate revocation lists, clearing..... | 1848 |
| certificate revocation lists, | |
| displaying..... | 1936 |
| certificate revocation lists, loading | |
| manually..... | 1867 |
| key pair, generating..... | 1870 |
| local certificates, clearing..... | 1849, 1850 |
| local certificates, displaying..... | 1938 |
| local certificates, loading manually..... | 1874 |
| local certificates, requesting | |
| manually..... | 1868, 1873 |
| local certificates, requesting online..... | 1864 |
| local certificates, requesting that CA | |
| install..... | 1871 |
| local certificates, requests, clearing..... | 1847 |
| local certificates, requests, | |
| displaying..... | 1934 |
| provided by Juniper Networks, adding..... | 2304 |
| signed certificate, obtaining..... | 2300 |
| unsigned certificate, obtaining..... | 2302 |
| cflowd statement | |
| usage guidelines..... | 898, 1078 |
| cgn-pic statement..... | 1326 |
| NAT..... | 75 |
| CGNAT | |
| ALGs..... | 141, 300 |
| chain-order statement | |
| nested applications..... | 1549 |
| CIR..... | 537 |
| cisco-interoperability statement..... | 1327 |
| usage guidelines..... | 547 |
| class statement..... | 1328 |
| clear ike security-associations command..... | 2296 |
| clear ipsec security-associations command..... | 2297 |
| clear passive-monitoring statistics | |
| command..... | 2164 |
| clear security pki ca-certificate command..... | 1846 |

| | | | |
|--|------|---|----------------------|
| clear security pki certificate-request
command..... | 1847 | clear services nat mappings app
command..... | 1841, 1842, 1844 |
| clear security pki crl command..... | 1848 | clear services nat mappings command..... | 1839 |
| clear security pki key-pair..... | 1849 | clear services rpm twamp server connection
command..... | 2168 |
| clear security pki local-certificate command..... | 1850 | clear services service-sets statistics packet-drops
command..... | 1852 |
| clear services accounting statistics inline-jflow
command..... | 2165 | clear services service-sets statistics syslog
command..... | 1853 |
| clear services application-aware-access-list
statistics command..... | 2083 | clear services stateful-firewall flows
command..... | 1854 |
| clear services application-identification
application-system-cache command..... | 2084 | clear services stateful-firewall sip-call
command..... | 1856 |
| clear services application-identification counter
command..... | 2085 | clear services stateful-firewall sip-register
command..... | 1859 |
| clear services cos statistics command..... | 1815 | clear services stateful-firewall statistics
command..... | 1862 |
| clear services crtp statistics command..... | 1816 | clear services video-monitoring mdi errors
command..... | 2169 |
| clear services dynamic-flow-capture
command..... | 2166 | clear services video-monitoring mdi statistics
command..... | 2170 |
| clear services flow-collector statistics
command..... | 2167 | clear-dont-fragment-bit statement
GRE tunnel..... | 1329 |
| clear services flows ip-action command..... | 2086 | IPsec..... | 1329 |
| clear services ids command..... | 1817 | usage guidelines..... | 426 |
| clear services ids destination-table
command..... | 1818 | service-set..... | 1330 |
| clear services ids pair-table command..... | 1819 | usage guidelines..... | 427, 434, 1208, 1216 |
| clear services ids source-table command..... | 1820 | clear-ike-sas-on-pic-restart statement..... | 1330 |
| clear services inline nat pool command..... | 1821 | usage guidelines..... | 390 |
| clear services inline nat statistics command..... | 1822 | clear-ipsec-sas-on-pic-restart statement..... | 1331 |
| clear services ipsec-vpn certificates
command..... | 1823 | usage guidelines..... | 390 |
| clear services ipsec-vpn ike security-associations
command..... | 1824 | client-list statement..... | 1645 |
| clear services ipsec-vpn ipsec security-associations
command..... | 1825 | close-timeout statement..... | 1269 |
| clear services ipsec-vpn ipsec statistics
command..... | 1826 | collector statement..... | 1645 |
| clear services l2tp destination command..... | 1827 | usage guidelines..... | 842 |
| clear services l2tp destination statistics
command..... | 1828 | collector-pic statement
usage guidelines..... | 843 |
| clear services l2tp multilink command..... | 1829 | command-name command..... | 2091 |
| clear services l2tp session command..... | 1830 | comments, in configuration statements..... | lx |
| clear services l2tp session statistics
command..... | 1832 | compression statement..... | 1331 |
| clear services l2tp tunnel command..... | 1834 | usage guidelines..... | 623, 624 |
| clear services l2tp tunnel statistics
command..... | 1836 | compression-device statement..... | 1332 |
| clear services local-policy-decision-function
statistics command..... | 2087 | usage guidelines..... | 626 |
| clear services nat flows command..... | 1838 | configuration
dynamic flow capture interface..... | 855 |
| | | flow-tap application..... | 865 |
| | | configuring NAT-PT with DNS application-level
gateways
example..... | 150 |

content destination
 dynamic flow capture, displaying.....2232

content destinations
 DFC.....847
 Junos Packet Vision.....860

content-destination statement.....1646
 usage guidelines.....850

context statement
 nested applications.....1550

control source
 DFC.....847

control source,
 dynamic flow capture, displaying.....2234

control-source statement.....1647
 usage guidelines.....851

conventions
 text and syntax.....lix

copy-dont-fragment-bit statement
 IPsec.....1332
 service-set.....1333

copy-tos-to-outer-ip-header statement.....1787
 usage guidelines.....1217

core-dump statement.....1648
 usage guidelines.....818

CoS
 action statements.....516
 applications.....515
 example configuration.....518
 for tunnels
 GRE TOS bits.....1217
 link services interfaces.....527, 562, 801
 match conditions.....515
 rules.....519
 scheduler map
 configuration example.....775

CoS services
 clear statistics.....1815
 mapping, displaying
 code point aliases to bit patterns.....1941

critical (system logging severity level).....21, 47, 645

CRTTP services
 flows, displaying.....1946
 output, displaying.....1944
 statistics, clearing.....1816

curly braces, in configuration statements.....lx

customer support.....lx
 contacting JTAC.....lx

D

data session identification
 APPID.....677

data statement.....1333
 usage guidelines.....517

data-fill statement.....1649

data-format statement.....1649
 usage guidelines.....841

data-size statement.....1650

dead peer detection (DPD) protocol.....427

delay buffer
 calculating.....532, 537
 shaping rate.....532, 537

delay-buffer-rate statement
 usage guidelines.....532

description statement
 IKE.....1334
 usage guidelines.....413
 IPsec.....1334
 usage guidelines.....418, 420

destination NAT
 configuring.....111, 183, 186

destination statement.....1600, 1651, 1788
 APPID
 usage guidelines.....676
 application identification rule.....1550
 encryption
 usage guidelines.....1251, 1260
 flow monitoring
 usage guidelines.....871, 1051
 link services
 usage guidelines.....721
 tunnel
 usage guidelines.....1213, 1241

destination-address statement
 AACL.....1551
 usage guidelines.....662
 CoS.....1335
 IDS.....1335
 usage guidelines.....357
 IPsec.....1336
 usage guidelines.....424
 NAT.....1336
 usage guidelines.....70
 stateful firewall.....1337
 usage guidelines.....332
 usage guidelines.....515

| | | |
|--|---------------|--|
| destination-address-range statement | | |
| ACL..... | 1551 | |
| usage guidelines..... | 662 | |
| IDS..... | 1337 | |
| usage guidelines..... | 357 | |
| NAT..... | 1338 | |
| usage guidelines..... | 70 | |
| stateful firewall..... | 1338 | |
| usage guidelines..... | 332 | |
| destination-interface statement | | |
| RPM..... | 1652 | |
| destination-ipv4-address (RFC 2544 | | |
| Benchmarking)..... | 1653 | |
| destination-mac-address (RFC2544 | | |
| Benchmarking)..... | 1653 | |
| destination-networks statement | | |
| tunnel..... | 1790 | |
| usage guidelines..... | 1245 | |
| destination-pool statement..... | 1339 | |
| usage guidelines..... | 71 | |
| destination-port range statement | | |
| NAT..... | 1341 | |
| destination-port statement | | |
| applications..... | 1267, 1315 | |
| RPM..... | 1340, 1654 | |
| usage guidelines..... | 309 | |
| destination-prefix statement..... | 1341, 1342 | |
| usage guidelines..... | 358 | |
| destination-prefix-ipv6 statement..... | 1342 | |
| usage guidelines..... | 358 | |
| destination-prefix-list statement | | |
| ACL..... | 1552 | |
| usage guidelines..... | 662 | |
| CoS..... | 1343 | |
| IDS..... | 1343 | |
| NAT..... | 1344 | |
| stateful firewall..... | 1344 | |
| usage guidelines..... | 332 | |
| destination-udp-port (RFC 2544 | | |
| Benchmarking)..... | 1655 | |
| destinations statement | | |
| flow collection..... | 1655 | |
| usage guidelines..... | 840 | |
| destined-port statement | | |
| NAT..... | 1345 | |
| deterministic-port-block-allocation | | |
| statement..... | 1346 | |
| DFC | | |
| architecture..... | 847 | |
| capture group..... | 849 | |
| control source configuration..... | 851 | |
| destination configuration..... | 850 | |
| example configuration..... | 855 | |
| interface configuration..... | 852 | |
| system logging..... | 853 | |
| threshold configuration..... | 854 | |
| dh-group statement..... | 1347 | |
| usage guidelines..... | 407 | |
| dial-options statement..... | 1348 | |
| interfaces | | |
| usage guidelines..... | 646 | |
| digital certificates See certificates | | |
| direction (RFC2544 Benchmarking)..... | 1656 | |
| direction statement..... | 1349 | |
| nested applications..... | 1552 | |
| usage guidelines..... | 386 | |
| disable statement | | |
| APPID | | |
| usage guidelines..... | 674, 680 | |
| application..... | 1553 | |
| application group..... | 1553 | |
| flow monitoring..... | 1657 | |
| port mapping..... | 1553 | |
| traffic sampling | | |
| usage guidelines..... | 874, 1054 | |
| disable-global-timeout-override statement..... | 1554 | |
| usage guidelines..... | 674 | |
| disable-mlppp-inner-ppp-pfc statement..... | 1601 | |
| usage guidelines..... | 730 | |
| discard accounting | | |
| usage guidelines..... | 883, 1063 | |
| dlci statement..... | 1601 | |
| usage guidelines..... | 777 | |
| DLCIs | | |
| multicast-capable connections..... | 778 | |
| point-to-point connections..... | 777 | |
| dnat-44 option | | |
| usage guidelines..... | 111, 183, 186 | |
| do-not-fragment statement | | |
| tunnel..... | 1791 | |
| usage guidelines..... | 1217 | |
| documentation | | |
| comments on..... | lxi | |
| download statement | | |
| APPID..... | 1554 | |
| usage guidelines..... | 686 | |

- drop-member-traffic statement
 - aggregated Multiservices.....1350
 - drop-timeout statement.....1602
 - usage guidelines.....730
 - ds-lite
 - subnet session limitation
 - configuring.....243
 - DS-Lite flows
 - statistics.....2035
 - ds-lite statement.....1351
 - usage guidelines.....227
 - dscp statement.....1352
 - usage guidelines.....516
 - dscp-code-point statement
 - RPM.....1658
 - DTCP.....847, 859
 - duplicates-dropped-periodicity statement.....1659
 - usage guidelines.....855
 - dynamic address-only source translation
 - configuring.....191
 - dynamic authentication.....456
 - dynamic flow capture *See* DFC
 - content destination, displaying.....2232
 - control source, displaying.....2234
 - statistics
 - clearing.....2166
 - displaying.....2236
 - dynamic flow capture interfaces
 - displaying.....2179
 - dynamic NAT
 - configuring.....191
 - dynamic route insertion.....457
 - dynamic rules.....456
 - dynamic security associations
 - usage guidelines.....390, 405
 - dynamic statement.....1352
 - usage guidelines.....390
 - Dynamic Tasking Control Protocol *See* DTCP
 - dynamic tunnels.....1792
 - destination.....1790
 - source.....1804
 - dynamic-flow-capture statement.....1660
 - dynamic-nat44 option
 - usage guidelines.....191
 - dynamic-tunnels statement.....1792
 - usage guidelines.....1245
- E**
- ecmp-alb statement.....1353
 - ei-mapping-timeout statement.....1354
 - emergency (system logging severity level).....21, 47, 644
 - enable flow collection mode.....844
 - enable-asymmetric-traffic-processing
 - statement.....1555
 - enable-heuristics statement.....1555
 - usage guidelines.....684
 - enable-rejoin statement
 - aggregated Multiservices.....1355
 - encapsulation statement.....1356
 - usage guidelines.....729
 - voice services
 - usage guidelines.....625
 - encrypted traffic identification
 - APPID.....684
 - encryption interface.....1251
 - applying inbound filter.....1257
 - example configuration.....1257
 - applying outbound filter.....1256
 - example configuration.....1255, 1256
 - configuring inbound filter.....1256
 - example configuration.....1257
 - configuring MTU.....1252
 - encryption interfaces
 - status information, displaying.....2309
 - encryption statement.....1357
 - usage guidelines.....389
 - encryption-algorithm statement
 - IKE.....1358
 - usage guidelines.....408
 - IPsec.....1358
 - usage guidelines.....418
 - engine-id statement
 - flow monitoring.....1661
 - engine-type statement.....1662
 - error (system logging severity level).....21, 47, 645
 - ES interfaces
 - example configuration.....1252
 - ES PIC
 - apply inbound filter.....1257
 - PIC redundancy.....1259
 - redundancy
 - example configuration.....1259
 - tunnel redundancy.....1260
 - es-options statement.....1793
 - usage guidelines.....1259
 - establish-tunnels statement.....1359

| | |
|---------------------------------------|----------------------|
| event policy | |
| all (tracing flag)..... | 50 |
| APPID..... | 688 |
| configuration (tracing flag)..... | 50 |
| database (tracing flag)..... | 50 |
| events (tracing flag)..... | 50 |
| policy (tracing flag)..... | 50 |
| export-format statement..... | 1663 |
| usage guidelines..... | 821 |
| extension-service statement..... | 1664 |
| F | |
| f-max-period statement..... | 1359 |
| usage guidelines..... | 623 |
| facility-override statement..... | 1271, 1360, 1361 |
| usage guidelines..... | 47 |
| family (RFC2544 Benchmarking)..... | 1666 |
| family statement | |
| aggregated Multiservices..... | 1361 |
| encryption..... | 1794 |
| usage guidelines..... | 1251 |
| flow monitoring | |
| usage guidelines..... | 871, 1051 |
| interfaces..... | 1362 |
| usage guidelines..... | 45 |
| link services..... | 1605 |
| usage guidelines..... | 721 |
| voice services..... | 1363 |
| file statement..... | 1669 |
| L-PDF statistics..... | 1556 |
| traffic sampling..... | 1668 |
| traffic sampling output | |
| usage guidelines..... | 876, 878, 1056, 1058 |
| file-specification statement | |
| usage guidelines..... | 841, 842 |
| filename statement..... | 1670 |
| filename-prefix statement..... | 1671 |
| usage guidelines..... | 842 |
| files | |
| logging information output file..... | 878, 1058 |
| traffic sampling output files..... | 876, 1056 |
| var/log/sampled file..... | 878, 1058 |
| var/tmp/sampled.pkts file..... | 876, 1056 |
| files statement..... | 1671 |
| usage guidelines..... | 876, 1056 |
| filter statement | |
| encryption..... | 1795 |
| usage guidelines..... | 1257 |
| flow monitoring..... | 1672 |
| usage guidelines..... | 871, 1051 |
| filters | |
| used with services..... | 31 |
| firewall filters | |
| actions..... | 871, 1051 |
| in traffic sampling..... | 871, 1051 |
| service filters..... | 39 |
| flow aggregation..... | 898, 1078 |
| multiple flow servers..... | 926, 1106 |
| source ID, IPFIX flows..... | 918, 1098 |
| template and option template ID..... | 921, 1101 |
| templates..... | 2208 |
| traffic sampling | |
| observation domain ID, version 9 | |
| | 918, 1098 |
| flow collector | |
| analyzer configuration..... | 841 |
| destination configuration..... | 840 |
| file format configuration..... | 841 |
| interface mapping..... | 842 |
| transfer log..... | 842 |
| flow collector interfaces | |
| status information, displaying..... | 2183 |
| flow collector services | |
| interface files, displaying..... | 2239 |
| packets received, displaying..... | 2241 |
| primary server, switching to..... | 2171 |
| secondary server, switching to..... | 2172 |
| statistics | |
| displaying..... | 2243 |
| dropped-packet, clearing..... | 1852, 1853 |
| interface, clearing..... | 2167 |
| test file, transferring..... | 2173 |
| flow limiting..... | 37 |
| Flow monitoring | |
| overview..... | 815, 829 |
| flow monitoring | |
| active | |
| aggregated flows, displaying..... | 2204 |
| CPU usage, displaying..... | 2230 |
| detailed information, displaying..... | 2218 |
| error statistics, displaying..... | 2209 |
| flow statistics, displaying..... | 2213 |
| memory statistics, displaying..... | 2223 |

| | |
|---|----------------------|
| packet size distribution, displaying..... | 2225 |
| PICs, displaying available..... | 2227 |
| example configuration | |
| multiple port | |
| mirroring..... | 940, 949, 1120, 1128 |
| next-hop groups..... | 940, 949, 1120, 1128 |
| inline | |
| flow statistics, clearing..... | 2165 |
| passive | |
| flow statistics, displaying..... | 2196 |
| memory and flow statistics, | |
| displaying..... | 2198 |
| status, displaying..... | 2200 |
| usage statistics, displaying..... | 2202 |
| redundancy..... | 826 |
| flow monitoring interfaces | |
| status information, displaying..... | 2189 |
| flow server | |
| replicating flows to multiple | |
| servers..... | 926, 1106 |
| flow-active-timeout statement..... | 1673 |
| usage guidelines..... | 821 |
| flow-collector statement..... | 1674 |
| usage guidelines..... | 839, 844 |
| flow-export-destination statement..... | 1675 |
| usage guidelines..... | 821 |
| flow-export-rate statement | |
| flow monitoring..... | 1675 |
| flow-inactive-timeout statement..... | 1676 |
| usage guidelines..... | 821 |
| flow-server statement | |
| flow monitoring..... | 1677 |
| flow-tap | |
| interface..... | 861 |
| permissions statement..... | 862 |
| RADIUS configuration..... | 862 |
| restrictions..... | 863 |
| security..... | 862 |
| flow-tap application | |
| example configuration..... | 865 |
| flow-tap statement..... | 1679 |
| flow-tap-dtcp statement..... | 862 |
| flows | |
| access-list..... | 2091 |
| list-flows..... | 2091 |
| font conventions..... | lix |
| force-entry statement..... | 1364 |
| usage guidelines..... | 358 |
| forwarding classes | |
| fragmentation..... | 527 |
| forwarding classes, displaying..... | 1941 |
| forwarding-class statement..... | 1364, 1365 |
| usage guidelines..... | 516, 527 |
| forwarding-options statement | |
| usage guidelines..... | 1624 |
| fragment-threshold statement | |
| link services..... | 1606 |
| usage guidelines..... | 732 |
| LSQ..... | 1366 |
| usage guidelines..... | 527 |
| voice services..... | 1367 |
| usage guidelines..... | 624 |
| fragmentation | |
| forwarding classes..... | 527 |
| GRE tunnels..... | 1216 |
| multiclass MLPPP..... | 562 |
| fragmentation and reassembly..... | 624, 773 |
| example configuration..... | 775 |
| fragmentation-map statement..... | 1367 |
| usage guidelines..... | 527 |
| fragmentation-maps statement..... | 1368 |
| usage guidelines..... | 527 |
| Frame Relay connections | |
| point-to-point connections..... | 777 |
| Frame Relay encapsulation | |
| multicast-capable connections..... | 778 |
| FRF.12..... | 624 |
| example configuration..... | 585 |
| LFI..... | 773 |
| LSQ..... | 582 |
| FRF.16..... | 571 |
| configuration example..... | 574 |
| from statement | |
| AACL..... | 1557 |
| usage guidelines..... | 661 |
| CoS..... | 1369 |
| HCM..... | 1371 |
| usage guidelines..... | 635 |
| IDS..... | 1370 |
| usage guidelines..... | 355, 357 |
| IPsec..... | 1371 |
| usage guidelines..... | 422, 424 |
| NAT..... | 1372 |
| usage guidelines..... | 70 |
| stateful firewall..... | 1373 |
| usage guidelines..... | 331, 332 |
| usage guidelines..... | 514 |

| | |
|----------------------------|---------------|
| ftp statement..... | 1374 |
| usage guidelines..... | 517, 840, 842 |
| FTP traffic, sampling..... | 880, 1060 |

G

| | |
|---|------|
| g-duplicates-dropped-periodicity statement..... | 1682 |
| usage guidelines..... | 855 |
| g-max-duplicates statement..... | 1683 |
| usage guidelines..... | 855 |
| GRE interfaces | |
| status information, displaying..... | 2315 |
| GRE tunnels | |
| fragmentation..... | 1216 |
| key number..... | 1215 |
| guaranteed rate..... | 537 |
| guaranteed-rate statement | |
| usage guidelines..... | 537 |

H

| | |
|-------------------------------------|-------------|
| hard-limit statement..... | 1683 |
| usage guidelines..... | 850 |
| hard-limit-target statement..... | 1684 |
| usage guidelines..... | 850 |
| hardware requirements..... | 3 |
| hardware-timestamp statement..... | 1684 |
| hello-interval statement | |
| L2TP..... | 1374 |
| usage guidelines..... | 643 |
| hello-timer statement | |
| link services..... | 1606 |
| usage guidelines..... | 727 |
| heuristics support | |
| APPID..... | 684 |
| hide-avps statement..... | 1375 |
| usage guidelines..... | 644 |
| high-availability-options statement | |
| aggregated Multiservices..... | 1376 |
| history-size statement..... | 1685 |
| usage guidelines..... | 971, 1149 |
| hold-time statement | |
| GRE tunnel interface..... | 1795 |
| host statement..... | 1272, 1378 |
| HCM..... | 1379 |
| L2TP..... | 1377 |
| usage guidelines..... | 20, 47, 644 |
| host-outbound statement..... | 1685 |
| hot-standby statement..... | 1379 |

I

| | |
|---|----------|
| ICMP | |
| ALGs, supported on the MS-MIC and MS-MPC..... | 322 |
| icmp-code statement..... | 1380 |
| usage guidelines..... | 307 |
| icmp-type statement..... | 1380 |
| usage guidelines..... | 307 |
| idle-timeout statement..... | 1558 |
| APPID | |
| usage guidelines..... | 674 |
| IDS | |
| action statements..... | 358 |
| applications..... | 357 |
| example configurations..... | 364 |
| match conditions..... | 357 |
| rules..... | 355 |
| IDS events | |
| clearing | |
| for a destination..... | 1818 |
| for interfaces and services..... | 1817 |
| for source addresses..... | 1820 |
| for source and destination pairs..... | 1819 |
| displaying..... | 1948 |
| ids-rule-sets statement | |
| usage guidelines..... | 36 |
| ids-rules statement..... | 1381 |
| usage guidelines..... | 36 |
| ignore-entry statement..... | 1364 |
| usage guidelines..... | 358 |
| ignore-errors statement..... | 1558 |
| usage guidelines..... | 677 |
| IKE..... | 372, 405 |
| adaptive services interfaces | |
| security associations, clearing..... | 1824 |
| security associations, displaying..... | 1961 |
| statistics, clearing..... | 1826 |
| authentication algorithm | |
| usage guidelines..... | 406 |
| authentication-method statement | |
| usage guidelines..... | 406 |
| DH (Diffie-Hellman) group | |
| usage guidelines..... | 407 |
| dynamic SAs..... | 405 |
| encryption services interfaces | |
| security associations, clearing..... | 2296 |
| security associations, displaying..... | 2305 |
| encryption-algorithm statement | |
| usage guidelines..... | 408 |

- lifetime
 - usage guidelines.....408
- mode statement
 - usage guidelines.....411
- policy.....409
 - example.....414
- policy statement
 - usage guidelines.....409
- pre-shared-key statement
 - usage guidelines.....411
- proposals statement
 - usage guidelines.....411
- supported software standards.....378
- version statement
 - usage guidelines.....411
- IKE profile
 - configuring access profile.....457
- IKE proposal
 - example configuration.....409
- IKE proposals
 - default.....460
- IKE security associations
 - clearing.....390
- ike statement.....1382
 - usage guidelines.....405
- ike-access-profile statement.....1383
 - usage guidelines.....432, 459
- in-service (RFC2544 Benchmarking).....1686
- inactivity-non-tcp-timeout statement.....1560
 - usage guidelines.....674
- inactivity-tcp-timeout statement.....1560
 - usage guidelines.....674
- inactivity-timeout statement.....1383
 - flow monitoring.....1272
 - RPM.....1687
 - usage guidelines.....19, 312
- index statement.....1559
 - APPID
 - usage guidelines.....674, 680
 - nested applications.....1559
- info (system logging severity level).....21, 48, 645
- initiate-dead-peer-detection statement.....1384
 - usage guidelines.....428
- inline flow monitoring
 - flow statistics, clearing.....2165
- inline FRF.15.....754
- inline FRF.16.....788
- inline LSQ services.....563
- inline MLPPP for WAN interfaces.....559
- inline NAT
 - statistics, displaying.....1956, 1957
- inline-jflow statement
 - flow monitoring.....1687
 - usage guidelines.....890, 894, 1070, 1074
- input statement
 - interfaces.....1384
 - usage guidelines.....31, 38
- input-interface-index statement.....1689
- input-packet-rate-threshold statement.....1689
 - usage guidelines.....854
- inside and outside interfaces.....34
- inside-service-interface statement
 - usage guidelines.....35
- instance statement
 - sampling.....1690
 - usage guidelines.....881, 1061
- interchassis LSQ failover.....545
- interface preservation.....551
- interface statement
 - encryption
 - usage guidelines.....1251
 - flow monitoring
 - usage guidelines.....934, 1114
 - flow-tap.....1692
 - usage guidelines.....861
- interface style service sets.....34
- interface-map statement.....1692
 - usage guidelines.....842
- interface-service statement.....1385
 - usage guidelines.....31
- interfaces
 - naming.....9
- interfaces statement
 - aggregated Multiservices.....1386
 - DFC.....1693
 - usage guidelines852
 - encryption.....1796
 - usage guidelines.....1251
 - flow monitoring
 - usage guidelines.....871, 1051
 - interfaces hierarchy.....1273
 - link services.....1607
 - tunnel
 - usage guidelines.....1213
 - video-monitoring.....1694
 - voice services.....1387
- interleave-fragments statement.....1607
 - usage guidelines.....773

| | |
|--|-----------|
| Internet Key Exchange See IKE | |
| intrachassis LSQ failover..... | 548 |
| intrusion detection | |
| example configurations..... | 364 |
| rule set..... | 363 |
| invalid SPI recovery | |
| enabling..... | 414 |
| IP addresses | |
| sampling traffic from single IP | |
| addresses..... | 879, 1059 |
| ip statement | |
| APPID | |
| usage guidelines..... | 676 |
| application identification..... | 1561 |
| ip-action | |
| clearing..... | 2086 |
| IP-over-IP interfaces | |
| status information, displaying..... | 2322 |
| ip-swap (RFC 2544 Benchmarking)..... | 1695 |
| IPsec | |
| action statements..... | 426 |
| authentication statement | |
| usage guidelines..... | 388 |
| authentication-algorithm statement | |
| usage guidelines..... | 416 |
| direction | |
| usage guidelines..... | 386 |
| dynamic authentication..... | 456 |
| dynamic endpoints for IPsec tunnels..... | 455 |
| dynamic endpoints interface | |
| configuration..... | 459 |
| dynamic rules..... | 456 |
| dynamic security associations | |
| usage guidelines..... | 390 |
| encryption | |
| usage guidelines..... | 389 |
| encryption-algorithm statement | |
| usage guidelines..... | 418 |
| ES PIC..... | 1251 |
| example configuration | |
| inbound traffic..... | 1257 |
| outbound traffic..... | 1255 |
| example policy configuration..... | 422 |
| IKE..... | 372 |
| lifetime of SA..... | 418 |
| lifetime-seconds statement..... | 418 |
| match conditions..... | 424 |
| minimum configurations | |
| dynamic SA | 384 |
| manual SA | 383 |
| overview..... | 371 |
| perfect-forward-secrecy statement | |
| usage guidelines..... | 421 |
| policy | |
| overview..... | 420 |
| policy statement | |
| usage guidelines..... | 420 |
| proposal statement | |
| usage guidelines..... | 415 |
| proposals statement | |
| usage guidelines..... | 421 |
| protocol statement (dynamic SA) | |
| usage guidelines..... | 419 |
| protocol statement (manual SA) | |
| usage guidelines..... | 387 |
| rule sets..... | 429 |
| security associations..... | 372 |
| security parameter index | |
| usage guidelines..... | 387 |
| service set dynamic endpoints | |
| configuration..... | 459 |
| Services SDK | |
| configuration..... | 507 |
| supported software standards..... | 378 |
| traffic..... | 1253 |
| IPsec proposals | |
| default..... | 460 |
| IPsec rules | |
| match directions..... | 424 |
| IPsec services | |
| adaptive services interfaces | |
| backup and primary, switching | |
| tunnels..... | 1876 |
| IKE security associations, clearing..... | 1824 |
| IKE security associations, displaying..... | 1961 |
| IPSec security associations, | |
| clearing..... | 1825 |
| IPSec security associations, | |
| displaying..... | 1965 |
| IPSec statistics, clearing..... | 1826 |
| IPSec statistics, displaying..... | 1969 |
| encryption services interfaces | |
| backup and primary, switching | |
| interfaces..... | 2299 |
| backup and primary, switching | |
| services..... | 2299 |

- certificate database, displaying.....2345
 - IKE security associations, clearing.....2296
 - IKE security associations,
 - displaying.....2305
 - IPSec security associations,
 - clearing.....2297
 - IPSec security associations,
 - displaying.....2350
 - redundancy information, displaying.....2348
 - ipsec statement.....1388
 - usage guidelines.....415
 - ipsec-inside-interface
 - usage guidelines.....456
 - ipsec-inside-interface statement.....1388
 - usage guidelines.....424
 - ipsec-interface-id statement
 - usage guidelines.....459
 - ipsec-sa statement
 - encryption.....1796
 - usage guidelines.....1251
 - ipsec-vpn-options statement.....1389
 - usage guidelines.....431
 - ipsec-vpn-rule-sets statement
 - usage guidelines.....36
 - ipsec-vpn-rules statement.....1389
 - usage guidelines.....36
 - IPv4
 - napt-44 option.....127
 - translation type
 - basic-nat-pt option.....143
 - basic-nat44 option.....93
 - basic-nat66 option.....99
 - IPv4 dynamic source translation
 - configuring.....127
 - IPv4 static source translation
 - AMS.....616
 - example.....616
 - ipv4-template statement.....1696
 - IPv6
 - napt-66 option.....131
 - transition
 - configured tunnel.....1211
 - IPv6 dynamic source translation
 - configuring.....131
 - ipv6-multicast-interfaces statement.....1390
 - IPv6-over-IPv4 tunnel
 - example configuration.....1211
 - standards supported.....1211
 - ipv6-template statement.....1697
- J**
- Junos Network Secure.....327
 - overview.....327
 - See also stateful firewall
 - Junos Packet Vision
 - application.....859
 - architecture.....860
- K**
- keepalive-time statement
 - GRE tunnel interface.....1797
 - key pair for digital certificate, generating.....2303
 - key statement
 - tunnel.....1798
 - usage guidelines.....1215
- L**
- L-PDF
 - best-effort application
 - identification.....658, 671, 693, 699
 - statistics
 - clearing.....2087
 - L2TP
 - access profile.....641, 642
 - attribute-value pairs.....644
 - example configuration.....648
 - redundancy.....648
 - timers.....643
 - L2TP LAC services
 - destination
 - clearing.....1827, 1828
 - L2TP services
 - multilink sessions
 - clearing.....1829
 - displaying.....1976
 - RADIUS information.....1980
 - session statistics
 - clearing.....1832
 - sessions
 - clearing.....1830
 - displaying.....1984
 - summary information, displaying.....1992
 - tunnel statistics, clearing.....1836
 - tunnels, clearing.....1834
 - tunnels, displaying.....1997
 - user information, displaying.....2003

| | |
|---|--------------------|
| L2TP statements | |
| LAC | |
| traceoptions..... | 1511 |
| LNS | |
| l2tp-access-profile..... | 1390 |
| local-gateway..... | 1395 |
| service-interface..... | 1466 |
| traceoptions..... | 1511 |
| l2tp-access-profile statement..... | 1390 |
| usage guidelines..... | 642 |
| l2tp-interface-id statement | |
| usage guidelines..... | 646 |
| l2tp-profile statement | |
| usage guidelines..... | 641 |
| label-position statement..... | 1697 |
| lawful intercept architecture..... | 860 |
| learn-sip-register statement..... | 1391 |
| LFI..... | 577, 582, 624, 773 |
| example configuration..... | 580, 585, 775 |
| lifetime-seconds statement | |
| IKE..... | 1391 |
| usage guidelines..... | 408 |
| IPsec..... | 1391 |
| usage guidelines..... | 418 |
| limiting flows per service set..... | 37 |
| link fragmentation and interleaving See LFI | |
| link PIC redundancy..... | 551 |
| link services interface | |
| multilink bundles See multilink bundles | |
| link services interfaces | |
| CoS components..... | 527, 562, 801 |
| example configuration..... | 803 |
| interleave fragments..... | 773 |
| example configuration..... | 775 |
| status information, displaying..... | 2118 |
| link services IQ interfaces..... | 580 |
| example configuration..... | 568, 574 |
| link state replication..... | 551 |
| link-layer overhead..... | 561 |
| status information, displaying..... | 1885, 2131 |
| link services protocols..... | 711 |
| link state replication | |
| LSQ PICs..... | 551 |
| link-layer overhead | |
| link services IQ interfaces..... | 561 |
| link-layer-overhead statement..... | 1392 |
| usage guidelines..... | 524, 529, 561 |
| list-flows..... | 2091 |
| See also list-statistics | |
| lmi-type statement..... | 1608 |
| usage guidelines..... | 728 |
| load balancing..... | 1925 |
| load-balance statement..... | 1392 |
| load-balancing-options statement | |
| aggregated Multiservices..... | 1393 |
| local-certificate statement..... | 1394 |
| usage guidelines..... | 412 |
| local-dump statement..... | 1698 |
| usage guidelines..... | 928, 1108 |
| local-gateway address statement | |
| usage guidelines..... | 642 |
| local-gateway statement..... | 1395 |
| usage guidelines..... | 431 |
| local-id statement..... | 1396 |
| usage guidelines..... | 413 |
| local-policy-decision-function statement..... | 1562 |
| log output | |
| adaptive services..... | 49 |
| APPID..... | 687 |
| traffic sampling..... | 878, 1058 |
| log-prefix statement..... | 1273, 1397 |
| L2TP..... | 1396 |
| usage guidelines..... | 20, 47, 644 |
| logging statement..... | 1397, 1398 |
| usage guidelines..... | 358 |
| logical interface scheduling..... | 1424 |
| logical interfaces | |
| multicast-capable connections..... | 778 |
| logical tunnel interfaces | |
| status information, displaying..... | 2326 |
| logical tunnels..... | 1221 |
| example configuration..... | 1222 |
| logical-system statement | |
| RPM..... | 1698 |
| usage guidelines..... | 971, 1149 |
| loopback tunnels..... | 1239 |
| LSQ | |
| CPU usage information, displaying..... | 1972 |
| LSQ bandwidth | |
| oversubscribing..... | 532 |
| LSQ failover | |
| interchassis..... | 545 |
| stateful intrachassis..... | 548 |
| stateless intrachassis..... | 548 |
| LSQ PICs..... | 551 |
| redundancy..... | 548 |
| lsq-failure-options statement..... | 1398 |
| usage guidelines..... | 546 |

M

- manual security association.....385
- manual statement.....1399
 - usage guidelines.....385
- manuals
 - comments on.....lxi
- many-to-one statement
 - aggregated Multiservices.....1400
- mapping-timeout statement.....1402
- match direction usage in service sets.....34
- match statement.....1699
- match-direction statement
 - AACL.....1563
 - usage guidelines.....662
 - CoS.....1402
 - IDS.....1403
 - usage guidelines.....355
 - IPsec.....1403
 - usage guidelines.....422
 - NAT.....1404
 - stateful firewall.....1404
 - usage guidelines.....332
 - usage guidelines.....515
- max-checked-bytes statement.....1564
 - APPID
 - usage guidelines.....683
- max-connection-duration statement.....1699
- max-drop-flows statement.....1405
- max-duplicates statement.....1700
 - usage guidelines.....855
- max-flows statement.....1406
 - usage guidelines.....37
- max-packets-per-second statement.....1701
 - usage guidelines.....873, 1053
- maximum-age statement.....1701
 - usage guidelines.....842
- maximum-connections statement.....1702
- maximum-connections-per-client
 - statement.....1703
- maximum-contexts statement.....1407
 - usage guidelines.....623
- maximum-packet-length statement.....1704
- maximum-send-window statement.....1407
 - usage guidelines.....643
- maximum-sessions statement.....1705
- maximum-sessions-per-connection
 - statement.....1706
- maximum-transactions statement
 - nested applications.....1564
- media delivery index
 - delay factor.....1043, 1191
 - media loss rate.....1043, 1191
 - media rate variation.....1043, 1191
- mediation devices
 - Junos Packet Vision.....860
- member statement
 - nested applications.....1565
- member-failure-options statement
 - aggregated Multiservices.....1408
- member-interface statement
 - aggregated Multiservices.....1410
- min-checked-bytes statement.....1565
 - APPID
 - usage guidelines.....683
- minimum links
 - link services interfaces.....733
 - multilink interfaces.....733
- minimum-links statement.....1608
 - usage guidelines.....733
- minimum-priority statement.....1706
 - usage guidelines.....851
- mlfr-uni-nni-bundle-options statement.....1609
 - usage guidelines.....725, 728
- mlfr-uni-nni-bundles-inline.....1412
- MLPPP.....565, 577
 - configuration example.....568
 - example configuration.....580
- MLPPP (Multilink Point-to-Point Protocol)
 - sample topology.....750, 755, 789
- mode (RFC 2544 Benchmarking).....1707
- mode statement.....1413
 - usage guidelines.....411
- monitoring statement.....1708
 - usage guidelines.....820
- moving-average-size statement.....1709
- MPLS
 - packets
 - passive flow monitoring.....832
- mpls-ipv4-template statement.....1709
- mpls-template statement.....1710
- mrru statement.....1610
 - usage guidelines.....734
- MS-MIC and MS-MPC
 - supported ALGs
 - identical support as uKernel.....322
- MS-MPC
 - configuration example
 - napt.....134

| | |
|--------------------------------------|---------------|
| mss statement..... | 1413 |
| usage guidelines..... | 358 |
| mtu statement..... | 1611 |
| multi-link-layer-2-inline..... | 1414 |
| multicast traffic | |
| AS PIC..... | 38 |
| multicast tunnel interfaces | |
| status information, displaying..... | 2331 |
| multicast tunnels..... | 1219 |
| multicast-capable connections | |
| Frame Relay encapsulation..... | 778 |
| multicast-dlci statement..... | 1611 |
| usage guidelines..... | 778 |
| multicast-only statement..... | 1798 |
| usage guidelines..... | 1219 |
| multiclass MLPPP | |
| fragmentation..... | 562 |
| multilink bundles | |
| fractional T1..... | 577 |
| example configuration..... | 580, 582, 585 |
| FRF.12..... | 582 |
| example configuration..... | 585 |
| MLPPP..... | 577 |
| example configuration..... | 580 |
| NxT1..... | 565, 571 |
| configuration example..... | 568, 574 |
| overview..... | 719 |
| sample topology..... | 750, 755, 789 |
| multilink interfaces | |
| minimum links..... | 733 |
| status information, displaying..... | 2155 |
| multilink-class statement..... | 1414 |
| usage guidelines..... | 562 |
| multilink-max-classes statement..... | 1415 |
| usage guidelines..... | 562 |
| multiservice-options statement..... | 1710 |
| MultiServices PIC | |
| hardware requirements..... | 25 |

N

| | |
|----------------------------|------|
| n391 statement..... | 1612 |
| usage guidelines..... | 728 |
| n392 statement..... | 1613 |
| usage guidelines..... | 728 |
| n393 statement..... | 1614 |
| usage guidelines..... | 728 |
| name-format statement..... | 1711 |
| usage guidelines..... | 841 |

NAPT

| | |
|---|----------|
| comparison of implementation methods..... | 126 |
| configuring..... | 127, 131 |
| IPv4..... | 127 |
| IPv6..... | 131 |
| port allocation | |
| round-robin..... | 118 |
| sequential..... | 118 |
| port block allocation..... | 121 |

nap

| | |
|----------------------------|-----|
| configuration example..... | 134 |
| nap-44 option | |
| usage guidelines..... | 127 |
| nap-66 option | |
| usage guidelines..... | 131 |
| nap-pt option | |
| example..... | 150 |

NAT

| | |
|--|----------|
| action statements..... | 71 |
| ALGs..... | 141, 300 |
| AMS..... | 605 |
| applications..... | 70 |
| destination NAT..... | 111, 183 |
| dynamic address-only source translation..... | 191 |
| dynamic NAT..... | 191 |
| dynamic source translation..... | 127, 131 |
| inline..... | 199 |
| configuring..... | 201 |
| inter-chassis high availability..... | 252 |
| ipv6-multicast-interfaces information, | |
| displaying..... | 2007 |
| load balancing, example..... | 616 |
| mapping information, address-pooling | |
| paired..... | 2009 |
| mapping information, displaying..... | 2009 |
| mapping information, endpoint-independent | |
| | 2009 |
| mapping information, pcp..... | 2009 |
| match conditions..... | 70 |
| NAPT | |
| configuring address pools..... | 117 |
| NAT-PT example..... | 150 |
| service sets..... | 75 |
| session logging..... | 209 |
| static destination address translation..... | 111, 183 |
| status information, displaying..... | 2014 |
| twice NAT | |
| description..... | 57 |

- nat
 - flows
 - clearing.....1838
 - mappings
 - clearing.....1839, 1841, 1842, 1844
- nat-options statement.....1415
- nat-rule-sets statement
 - usage guidelines.....36
- nat-rules statement.....1416
 - usage guidelines.....36
- nested-application statement
 - APPID.....1566
 - usage guidelines.....681
- nested-application-settings statement
 - APPID.....1567
- network address translation
 - configuration example
 - napt.....134
- network address translation See NAT
- networks
 - sample LFI and multilink bundle
 - topology.....750, 755, 789
 - sample multilink bundle and LFI
 - topology.....750, 755, 789
- next-hop group for port mirroring.....948
- next-hop groups.....931, 1111
- next-hop statement.....1712
 - next-hop groups
 - usage guidelines.....934, 1114
 - usage guidelines.....931, 1111
- next-hop style service sets.....35
- next-hop-group statement
 - forwarding-options.....1713
 - usage guidelines.....931, 934, 1111, 1114
- next-hop-service statement.....1274, 1417
 - usage guidelines.....33
- no-anti-replay statement.....1418
 - usage guidelines.....428, 433
- no-application-identification statement.....1567
 - APPID
 - usage guidelines.....683
- no-application-system-cache statement.....1568
 - APPID
 - usage guidelines.....683
- no-clear-application-system-cache
 - statement.....1568
 - APPID
 - usage guidelines.....683
- no-core-dump statement.....1648
 - usage guidelines.....818
- no-filter-check statement.....1713
 - usage guidelines.....931, 1111
- no-fragmentation statement.....1419
 - usage guidelines.....527
- no-ipsec-tunnel-in-traceroute statement.....1419
 - usage guidelines.....436
- no-local-dump statement.....1698
 - usage guidelines.....928, 1108
- no-nested-application statement.....1569
 - usage guidelines.....682
- no-per-unit-scheduler statement.....1420
- no-protocol-method statement.....1569
 - APPID
 - usage guidelines.....683
- no-remote-trace statement
 - flow monitoring.....1714
- no-signature-based statement.....1570
 - APPID
 - usage guidelines.....683
- no-stamp statement.....1758
 - usage guidelines.....876, 1056
- no-syslog statement
 - DFC.....1714
 - flow monitoring.....1758
 - usage guidelines.....853
- no-termination-request statement.....1420
 - usage guidelines.....546
- no-translation statement.....1421
 - usage guidelines.....71
- no-world-readable statement
 - flow monitoring.....1783
 - usage guidelines.....876, 1056
- notice (system logging severity level).....21, 47, 645
- notification-targets statement.....1715
 - usage guidelines.....851
- NxT1 bundles
 - FRF.16.....571
 - configuration example.....574
 - MLPPP.....565
 - configuration example.....568
- O**
 - observation-domain-id statement.....1716
 - offloading flows
 - configuring.....827
 - one-way-hardware-timestamp statement.....1717
 - usage guidelines.....966, 1144

| | |
|---|--------------------|
| open-timeout statement..... | 1275 |
| usage guidelines..... | 19 |
| option-refresh-rate statement..... | 1718 |
| options-template-id statement..... | 1719 |
| order statement..... | 1570 |
| APPID | |
| usage guidelines..... | 676 |
| output files | |
| logging information output file..... | 878, 1058 |
| traffic sampling output files..... | 876, 1056 |
| output statement..... | 1421 |
| discard accounting..... | 1720 |
| flow monitoring..... | 1721 |
| port mirroring..... | 1722 |
| sampling..... | 1723 |
| usage guidelines..... | 31, 38 |
| output-interface-index statement..... | 1724 |
| outside-service-interface statement | |
| usage guidelines..... | 35 |
| overload-pool statement..... | 1422 |
| usage guidelines..... | 71 |
| overload-prefix statement..... | 1422 |
| usage guidelines..... | 71 |
| oversubscription..... | 532 |
| P | |
| packet size distribution, displaying..... | 2225 |
| packet-based IPsec..... | 424 |
| parentheses, in syntax descriptions..... | lx |
| passive flow monitoring..... | 815, 829 |
| error statistics, displaying..... | 2194 |
| flow statistics, displaying..... | 2196 |
| memory statistics, displaying..... | 2198 |
| MPLS packets..... | 832 |
| PICs, displaying available..... | 2200 |
| statistics, clearing..... | 2164 |
| usage statistics, displaying..... | 2202 |
| passive-mode-tunneling statement..... | 1423 |
| usage guidelines..... | 435 |
| passive-monitor-mode statement..... | 1724 |
| usage guidelines..... | 830 |
| password statement | |
| usage guidelines..... | 840, 842 |
| pattern statement | |
| nested applications..... | 1571 |
| peer-as-billing-template statement..... | 1726 |
| peer-unit statement | |
| tunnel..... | 1799 |
| usage guidelines..... | 1221 |
| per-unit scheduling..... | 1424 |
| per-unit-scheduler statement..... | 1424 |
| usage guidelines..... | 532, 537, 565, 571 |
| perfect-forward-secrecy statement..... | 1425 |
| usage guidelines..... | 421 |
| pgcp-rules statement | |
| service-set..... | 1426 |
| PIC types for services..... | 3 |
| pic-memory-threshold statement..... | 1726 |
| usage guidelines..... | 854 |
| PICs | |
| active flow monitoring | |
| available PICs, displaying..... | 2227 |
| CPU usage, displaying..... | 2230 |
| PIM | |
| tunnels..... | 1237 |
| PIM de-encapsulation interfaces | |
| status information, displaying..... | 2336 |
| PIM encapsulation interfaces | |
| status information, displaying..... | 2336 |
| ping | |
| ALGs, supported on the MS-MIC and | |
| MS-MPC..... | 322 |
| PIR..... | 532 |
| PKI See certificates, PKI | |
| platforms, supported..... | 7 |
| point-to-point connections | |
| Frame Relay encapsulation..... | 777 |
| policy statement | |
| IKE..... | 1427 |
| usage guidelines..... | 409 |
| IPsec..... | 1428 |
| usage guidelines..... | 420 |
| policy-decision-statistics-profile statement..... | 1572 |
| pool statement..... | 1429 |
| pop-all-labels statement..... | 1727 |
| usage guidelines..... | 833 |
| port block allocation..... | 121 |
| deterministic..... | 122 |
| algorithms..... | 122 |
| configuring..... | 180 |
| interim syslog messages..... | 209 |
| secured..... | 121, 177 |
| configuring..... | 178 |
| Port Control Protocol | |
| Configuring..... | 167 |
| Configuring a Service Set to Apply PCP..... | 169 |
| Configuring PCP Server Options..... | 167, 168 |

- port forwarding
 - configuring.....186
 - dnat-44.....183
 - static destination address translation.....183
 - without destination address translation.....186
- port forwarding without static destination address translation
 - configuring.....186
- port mirroring.....931, 1111
 - disabling.....1657
 - displaying.....2177
- port statement
 - cflowd
 - usage guidelines.....898, 1078
 - flow monitoring.....1728
 - NAT.....1430
 - RPM.....1728
 - TWAMP.....1729
 - voice services.....1432
 - usage guidelines.....623
- port-forwarding
 - example.....187
- port-forwarding statement
 - destined-port statement.....1345
 - NAT.....1433
 - translated-port statement.....1518
- port-forwarding-mappings statement.....1433
- port-mapping statement.....1573
- port-mirroring statement
 - usage guidelines.....931, 1111
- port-range statement.....1573
 - APPID
 - usage guidelines.....676
- ports-per-session statement.....1434
- post-service-filter statement.....1434
 - usage guidelines.....31
- ppp-access-profile statement.....1435
 - usage guidelines.....642
- ppp-profile statement
 - usage guidelines.....641
- pre-rewrite-tos statement.....1729
 - usage guidelines.....875, 1055
- pre-shared-key statement.....1435
 - usage guidelines.....411
- preserve-interface statement.....1436
 - usage guidelines.....551
- primary statement
 - link services.....1437
 - usage guidelines.....549
 - services PIC.....1436
 - usage guidelines.....41
- probe statement
 - RPM.....1730
- probe-count statement.....1731
- probe-interval statement.....1732
- probe-limit statement.....1732
- probe-server statement.....1733
- probe-type statement.....1734
- procedural overview.....15
- profile statement
 - APPID
 - usage guidelines.....679
 - application identification.....1574
- proposal statement
 - IKE.....1437
 - usage guidelines.....405
 - IPsec.....1438
 - usage guidelines.....415
- proposals statement
 - IKE.....1438
 - usage guidelines.....411
 - IPsec.....1438
 - usage guidelines.....421
- protocol statement
 - applications.....1439
 - usage guidelines.....306
 - IPsec.....1440
 - usage guidelines.....387, 419
 - nested applications.....1574
- ptsp-rule-sets statement
 - usage guidelines.....36
- ptsp-rules statement.....1440
 - usage guidelines.....36
- Q**
 - queues statement.....1441
 - usage guidelines.....623
- R**
 - RADIUS information
 - displaying.....1980
 - RADIUS servers
 - configuration example.....616
 - random-allocation statement.....1430

| | | | |
|---|----------------------|--|------|
| rate statement..... | 1735 | request security certificate (unsigned) | |
| usage guidelines..... | 873, 931, 1053, 1111 | command..... | 2302 |
| reassemble-packets statement..... | 1799 | request security key-pair command..... | 2303 |
| usage guidelines..... | 1217 | request security pki ca-certificate enroll | |
| receive-options-packets statement..... | 1735 | command..... | 1864 |
| usage guidelines..... | 830 | request security pki ca-certificate load | |
| receive-ttl-exceeded statement..... | 1736 | command..... | 1865 |
| usage guidelines..... | 830 | request security pki ca-certificate verify | |
| receive-window statement..... | 1441 | command..... | 1866 |
| usage guidelines..... | 643 | request security pki crt load command..... | 1867 |
| red-differential-delay statement..... | 1614 | request security pki generate-certificate-request | |
| usage guidelines..... | 727 | command..... | 1868 |
| redistribute-all-traffic statement | | request security pki generate-key-pair | |
| aggregated Multiservices..... | 1442 | command..... | 1870 |
| redundancy | | request security pki local-certificate enroll | |
| AS PIC..... | 41 | command..... | 1871 |
| flow monitoring..... | 826 | request security pki local-certificate | |
| L2TP..... | 648 | generate-self-signed command..... | 1873 |
| redundancy group | | request security pki local-certificate load | |
| logical interfaces in..... | 1800 | command..... | 1874 |
| redundancy-group | | request security pki local-certificate verify | |
| logical tunnel..... | 1801 | command..... | 1875 |
| redundancy-options statement..... | 1443 | request services application-identification | |
| usage guidelines..... | 41 | application command..... | 2088 |
| redundant adaptive services interfaces | | request services application-identification group | |
| reverting to the primary interface..... | 1863 | command..... | 2090 |
| status information, displaying..... | 1909 | request services flow-collector change-destination | |
| switching to the secondary interface..... | 1863 | primary interface command..... | 2171 |
| redundant link services IQ interfaces | | request services flow-collector change-destination | |
| status information, displaying..... | 1911 | secondary interface command..... | 2172 |
| redundant logical tunnels | | request services flow-collector test-file-transfer | |
| overview..... | 1224 | command..... | 2173 |
| redundant logical tunnels | | request services ipsec-vpn ipsec switch tunnel | |
| configuring..... | 1226 | command..... | 1876 |
| reflect-mode (RFC2544 Benchmarking)..... | 1737 | request system certificate add command..... | 2304 |
| reflexive reverse statement..... | 1444 | request-url statement..... | 1446 |
| usage guidelines..... | 518 | required-depth statement..... | 1738 |
| rejoin-timeout statement | | usage guidelines..... | 833 |
| aggregated Multiservices..... | 1445 | respond-bad-ip statement | |
| remote-gateway statement..... | 1445 | usage guidelines..... | 414 |
| usage guidelines..... | 427 | retransmit-interval statement..... | 1447 |
| remote-id statement..... | 1446 | usage guidelines..... | 643 |
| usage guidelines..... | 413 | retry statement..... | 1739 |
| request interface (revert switchover) (Adaptive | | usage guidelines..... | 843 |
| Services) command..... | 1863 | retry-delay statement..... | 1739 |
| request ipsec switch command..... | 2299 | usage guidelines..... | 843 |
| request security certificate (signed) | | | |
| command..... | 2300 | | |

- RFC 2544 benchmarking test, RPM service
 - configuring.....989, 1163
 - example, configuring for Layer 3 IPv4 services.....993, 1168
 - example, configuring for NNI of Ethernet pseudowires.....1008, 1183
 - example, configuring for UNI of Ethernet pseudowires.....1001, 1175
 - layer 2 overview.....986
 - statistical details of a specific test ID, displaying.....2264
 - statistical details of a test type, displaying.....2259
 - test name, configuring.....989, 1163
 - test profile, configuring.....989, 1163
- RFC 2890.....1215
- RFC2544 benchmarking test, RPM service
 - overview.....983, 1161
- route-record statement
 - usage guidelines.....898, 1078
- routing-instance statement
 - RPM.....1741
 - tunnel.....1802
 - usage guidelines.....1241
- routing-instances statement.....1802
 - RPM.....1742
 - usage guidelines.....973, 1151
- routing-options statement.....1803
- rpc-program-number statement.....1447
 - usage guidelines.....320
- RPM.....957, 1135
 - example configuration.....977, 1155
- RPM services
 - benchmark test, performing.....2293
 - benchmarking test
 - configuring.....989, 1163
 - example, configuring for Layer 2 Reflection, ELAN, Bridge.....1016
 - example, configuring for Layer 3 IPv4 services.....993, 1168
 - example, configuring for NNI of Ethernet pseudowires.....1008, 1183
 - example, configuring for UNI of Ethernet pseudowires.....1001, 1175
 - layer 2 overview.....986
 - overview.....983, 1161
 - reflector commands.....988
 - benchmarking test results
 - displaying by test state.....2259
 - test ID, displaying.....2264
 - test type, displaying.....2259
 - displaying information of an RFC 2544 benchmarking test for a particular test type.....2259
 - displaying information of an RFC 2544 benchmarking test for a specific test ID.....2264
 - probe results
 - history, displaying.....2250
 - recent, displaying.....2253
 - protocols and ports, displaying.....2249
- rpm statement.....1742, 1743
- RPM statements
 - traceoptions.....1769
- RPM TWAMP server
 - connections, clearing.....2168
 - connections, displaying.....2281
 - sessions, displaying.....2283
- rtp statement.....1448
 - usage guidelines.....623
- rule statement
 - AACL.....1575
 - usage guidelines.....661
 - APPID
 - usage guidelines.....676
 - CoS.....1449
 - IDS.....1450
 - usage guidelines.....355
 - IPsec.....1452
 - usage guidelines.....422
 - NAT.....1454
 - software.....219, 1456
 - stateful firewall.....1455
 - usage guidelines.....331
 - usage guidelines.....514
- rule-set statement.....1456
 - AACL.....1577
 - usage guidelines.....666
 - APPID
 - usage guidelines.....676
 - application identification.....1578
 - CoS
 - usage guidelines.....519
 - IDS.....1457
 - usage guidelines.....363

| | |
|---------------------------|----------------------|
| IPsec..... | 1457 |
| usage guidelines..... | 429 |
| NAT..... | 1458, 1527 |
| software..... | 1459 |
| stateful firewall..... | 1276, 1458 |
| usage guidelines..... | 335 |
| run-length statement..... | 1745 |
| usage guidelines..... | 873, 931, 1053, 1111 |

S

| | |
|---|------------|
| sample (firewall filter action)..... | 871, 1051 |
| sample-once statement | |
| flow monitoring..... | 1745 |
| usage guidelines..... | 874, 1054 |
| sampled file..... | 878, 1058 |
| sampled.pkts file..... | 876, 1056 |
| sampling | |
| logical interface..... | 873, 1053 |
| monitoring interface..... | 818 |
| next-hop-groups, displaying..... | 2174 |
| port-mirroring instances, displaying..... | 2177 |
| sampling rate..... | 873, 1053 |
| sampling statement..... | 1746, 1748 |
| usage guidelines..... | 871, 1051 |
| scheduler map | |
| CoS | |
| configuration example..... | 775 |
| secondary statement | |
| link services..... | 1460 |
| usage guidelines..... | 549 |
| services PIC..... | 1459 |
| usage guidelines..... | 41 |
| secure-nat-mapping statement..... | 1460 |
| secured-port-block-allocation statement..... | 1461 |
| security associations | |
| clearing..... | 390 |
| configuring..... | 385 |
| security certificate See certificates | |
| senable-asymmetric-traffic-processing statement | |
| usage guidelines..... | 685 |
| send cflowd records to flow collector..... | 843 |
| server (PCP) statement..... | 1463 |
| server statement..... | 1749 |
| server-inactivity-timeout statement..... | 1749 |
| service filters..... | 39 |
| service interface configuration..... | 31 |
| service packages..... | 11 |
| service rules configuration..... | 36 |

| | |
|--|------------------------------|
| service sets | |
| example configuration..... | 41 |
| overview..... | 26 |
| service statement..... | 1464 |
| usage guidelines..... | 38 |
| service-domain statement..... | 1465 |
| usage guidelines..... | 33 |
| service-filter statement..... | 1465 |
| firewall | |
| usage guidelines..... | 39 |
| interfaces | |
| usage guidelines..... | 31 |
| service-interface statement..... | 1466 |
| usage guidelines..... | 31, 642 |
| service-port statement..... | 1750 |
| usage guidelines..... | 851 |
| service-set statement..... | 1276, 1277, 1467, 1468, 1580 |
| NAT..... | 75 |
| usage guidelines..... | 38 |
| service-set statements | |
| Adaptive Services interfaces | |
| service-interface..... | 1466 |
| service-type (RFC2544 Benchmarking)..... | 1750 |
| services configuration overview..... | 15 |
| services PICs..... | 3 |
| services sets | |
| CPU usage, displaying..... | 2022 |
| dropped packet statistics | |
| clearing..... | 1852 |
| displaying..... | 2026 |
| memory usage, displaying..... | 2024 |
| summary information, displaying..... | 2032 |
| syslog statistics | |
| clearing..... | 1853 |
| displaying..... | 2028 |
| services statement | |
| AACL | |
| usage guidelines..... | 1536 |
| APPID | |
| usage guidelines..... | 1295, 1537 |
| CoS..... | 1279 |
| DFC | |
| usage guidelines..... | 849 |
| dynamic-flow-control | |
| usage guidelines..... | 1628 |
| flow control | |
| usage guidelines..... | 1629 |
| flow-tap | |
| usage guidelines..... | 1631 |

| | | | |
|---|---------------|--|------|
| IDS..... | 1279 | show interfaces (Logical Tunnel) command..... | 2326 |
| interfaces..... | 1281 | show interfaces (Multicast Tunnel) | |
| usage guidelines..... | 20 | command..... | 2331 |
| IPsec..... | 1280 | show interfaces (Multilink Services) | |
| L2TP | | command..... | 2155 |
| usage guidelines..... | 644 | show interfaces (PIM) command..... | 2336 |
| NAT..... | 1281 | show interfaces (Redundant Adaptive Services) | |
| rpm | | command..... | 1909 |
| usage guidelines..... | 1631 | show interfaces (Redundant Link Services IQ) | |
| RPM..... | 1751 | command..... | 1911 |
| service sets | | show interfaces (Virtual Loopback Tunnel) | |
| usage guidelines..... | 47 | command..... | 2340 |
| stateful firewall..... | 1283 | show interfaces command..... | 1925 |
| services-options statement..... | 1285 | show interfaces redundancy command..... | 1928 |
| usage guidelines..... | 19, 20 | show ipsec certificates command..... | 2345 |
| session logging..... | 209 | show ipsec redundancy command..... | 2348 |
| session-limit statement..... | 1471 | show ipsec security-associations command..... | 2350 |
| usage guidelines..... | 358 | show passive-monitoring error command..... | 2194 |
| session-timeout statement..... | 1583 | show passive-monitoring flow command..... | 2196 |
| usage guidelines..... | 674 | show passive-monitoring memory | |
| set-dont-fragment-bit statement | | command..... | 2198 |
| IPsec..... | 1472 | show passive-monitoring status command..... | 2200 |
| service-set..... | 1472 | show passive-monitoring usage command..... | 2202 |
| shaping-rate statement | | show security pki ca-certificate command..... | 1930 |
| usage guidelines..... | 525, 532, 537 | show security pki certificate-request | |
| shared-key statement..... | 1751 | command..... | 1934 |
| usage guidelines..... | 851 | show security pki crt command..... | 1936 |
| short-sequence statement..... | 1615 | show security pki local-certificate command..... | 1938 |
| usage guidelines..... | 735 | show services accounting aggregation | |
| show application-identification | | command..... | 2204 |
| application-system-cache command..... | 2094 | show services accounting aggregation template | |
| show forwarding-options next-hop-group | | command..... | 2208 |
| command..... | 2174 | show services accounting errors command..... | 2209 |
| show forwarding-options port-mirroring | | show services accounting flow command..... | 2213 |
| command..... | 2177 | show services accounting flow-detail | |
| show ike security-associations command..... | 2305 | command..... | 2218 |
| show interfaces (Adaptive Services) | | show services accounting memory | |
| command..... | 1877 | command..... | 2223 |
| show interfaces (Dynamic Flow Capture) | | show services accounting packet-size-distribution | |
| command..... | 2179 | command..... | 2225 |
| show interfaces (Encryption) command..... | 2309 | show services accounting status command..... | 2227 |
| show interfaces (Flow Collector) command..... | 2183 | show services accounting usage command..... | 2230 |
| show interfaces (Flow Monitoring) | | show services application-aware-access-list | |
| command..... | 2189 | statistics command..... | 2101 |
| show interfaces (GRE) command..... | 2315 | show services application-identification application | |
| show interfaces (IP-over-IP) command..... | 2322 | command..... | 2103 |
| show interfaces (Link Services IQ) | | show services application-identification counter | |
| command..... | 1885, 2131 | command..... | 2096 |
| show interfaces (Link Services) command..... | 2118 | | |

| | | | |
|---|------|---|------|
| show services application-identification group
command..... | 2099 | show services rpm history-results command..... | 2250 |
| show services application-identification
version..... | 2106 | show services rpm probe-results command..... | 2253 |
| show services cos statistics command..... | 1941 | show services rpm rfc2544-benchmarking
command..... | 2259 |
| show services crtp command..... | 1944 | show services rpm rfc2544-benchmarking test-id
command..... | 2264 |
| show services crtp flows command..... | 1946 | show services rpm twamp server connection
command..... | 2281 |
| show services dynamic-flow-capture
content-destination command..... | 2232 | show services rpm twamp server session
command..... | 2283 |
| show services dynamic-flow-capture control-source
command..... | 2234 | show services service-sets cpu-usage
command..... | 2022 |
| show services dynamic-flow-capture statistics
command..... | 2236 | show services service-sets memory-usage
command..... | 2024 |
| show services flow-collector file interface
command..... | 2239 | show services service-sets statistics packet-drops
command..... | 2026 |
| show services flow-collector input interface
command..... | 2241 | show services service-sets statistics syslog
command..... | 2028 |
| show services flow-collector interface
command..... | 2243 | show services service-sets statistics tcp-mss
command..... | 2031 |
| show services flows command..... | 2107 | show services service-sets summary
command..... | 2032 |
| show services ids command..... | 1948 | show services softwire command..... | 2034 |
| show services inline nat pool command..... | 1956 | show services softwire flows command..... | 2035 |
| show services inline nat statistics command..... | 1957 | show services softwire statistics command..... | 2038 |
| show services ipsec-vpn certificates
command..... | 1958 | show services stateful-firewall conversations
command..... | 2044 |
| show services ipsec-vpn ike security-associations
command..... | 1961 | show services stateful-firewall flow-analysis
command..... | 2048 |
| show services ipsec-vpn ipsec security-associations
command..... | 1965 | show services stateful-firewall flows
command..... | 2052 |
| show services ipsec-vpn ipsec statistics
command..... | 1969 | show services stateful-firewall sip-call
command..... | 2058 |
| show services l2tp multilink command..... | 1976 | show services stateful-firewall sip-register
command..... | 2063 |
| show services l2tp radius command..... | 1980 | show services stateful-firewall statistics
application-protocol sip command..... | 2076 |
| show services l2tp session command..... | 1984 | show services stateful-firewall statistics
command..... | 2067 |
| show services l2tp summary command..... | 1992 | show services stateful-firewall subscriber analysis
command..... | 2079 |
| show services l2tp tunnel command..... | 1997 | show services video-monitoring mdi errors
command..... | 2285 |
| show services l2tp user command..... | 2003 | show services video-monitoring mdi flows
command..... | 2287 |
| show services link-services cpu-usage
command..... | 1972 | show services video-monitoring mdi stats
command..... | 2291 |
| show services local-policy-decision-function flows
command..... | 2114 | show system certificate command..... | 2353 |
| show services local-policy-decision-function
statistics command..... | 2116 | show system statistics command..... | 2304 |
| show services nat ipv6-multicast-interfaces
command..... | 2007 | | |
| show services nat mappings command..... | 2009 | | |
| show services nat pool command..... | 2014 | | |
| show services pcp statistics command..... | 2019 | | |
| show services rpm active-servers command..... | 2249 | | |

-
- signature statement
 - nested applications.....1584
 - signature-method-all-ports statement.....1584
 - AAPID
 - usage guidelines.....683
 - sip statement.....1474
 - usage guidelines.....517
 - sip-call-hold-timeout statement.....1473
 - size statement.....1752
 - usage guidelines.....878, 1058
 - snmp-command statement.....1474
 - usage guidelines.....320
 - soft-limit statement.....1753
 - usage guidelines.....850
 - soft-limit-clear statement.....1753
 - usage guidelines.....850
 - software flows
 - statistics.....2035
 - software-concentrator statement.....1475
 - software-rules statement.....1476
 - usage guidelines.....36
 - SONET interfaces
 - sampling SONET interfaces.....878, 1058
 - source statement
 - APPID
 - usage guidelines.....676
 - application identification rule.....1585
 - encryption.....1803
 - tunnel.....1804
 - usage guidelines.....1260
 - source-address statement
 - AACL.....1585
 - usage guidelines.....662
 - CoS.....1477
 - flow monitoring.....1754
 - usage guidelines.....820
 - IDS.....1478
 - usage guidelines.....357
 - IPsec.....1478
 - usage guidelines.....424
 - NAT.....1479
 - usage guidelines.....70
 - RPM.....1755
 - service-set system log.....1477
 - stateful firewall.....1479
 - usage guidelines.....332
 - tunnel.....1804
 - tunnel services
 - usage guidelines.....1245
 - usage guidelines.....515
 - source-address-range statement
 - AACL.....1586
 - usage guidelines.....662
 - IDS.....1480
 - usage guidelines.....357
 - NAT.....1480
 - usage guidelines.....70
 - stateful firewall.....1481
 - usage guidelines.....332
 - source-addresses statement
 - DFC.....1755
 - usage guidelines.....851
 - source-id statement.....1756
 - source-ipv4-address (RFC 2544 Benchmarking).....1756
 - source-mac-address (RFC2544 Benchmarking).....1757
 - source-pool statement.....1481
 - usage guidelines.....71
 - source-port statement
 - RPM.....1482
 - usage guidelines.....309
 - source-prefix statement.....1482, 1483
 - usage guidelines.....358
 - source-prefix-ipv6 statement.....1483
 - usage guidelines.....358
 - source-prefix-list statement
 - AACL.....1586
 - usage guidelines.....662
 - CoS.....1484
 - IDS.....1484, 1587
 - NAT.....1485
 - stateful firewall.....1485
 - usage guidelines.....332
 - source-udp-port (RFC 2544 Benchmarking).....1757
 - spi statement.....1486
 - usage guidelines.....387
 - stamp option.....877, 1057
 - stamp statement.....1758
 - usage guidelines.....876, 1056
 - stateful firewall
 - action statements.....334
 - anomalies.....328
 - applications.....332
 - conversations
 - displaying.....2044

| | | | |
|--|-----------|--|------------|
| example configuration..... | 335 | dynamic flow capture | |
| flow analysis | | clearing..... | 2166 |
| displaying..... | 2048 | displaying..... | 2236 |
| flows | | L-PDF | |
| clearing..... | 1854 | clearing..... | 2087 |
| displaying..... | 2052 | statistics statement | |
| match conditions..... | 332 | L-PDF..... | 1579 |
| overview..... | 327 | subnet session limitation | |
| rules..... | 335 | ds-lite | |
| SIP call information | | configuring..... | 243 |
| clearing..... | 1856 | support, technical See technical support | |
| displaying..... | 2058 | support-uni-directional-traffic statement..... | 1579 |
| SIP register information | | usage guidelines..... | 685 |
| clearing..... | 1859 | syn-cookie statement..... | 1492 |
| displaying..... | 2063 | usage guidelines..... | 358 |
| SIP statistics | | syntax conventions..... | lix |
| displaying..... | 2076 | syslog statement..... | 1487 |
| statistics | | flow monitoring..... | 1758 |
| clearing..... | 1862 | IDS..... | 1487 |
| displaying..... | 2067 | usage guidelines..... | 358 |
| subscriber analysis | | interfaces..... | 1286 |
| displaying..... | 2079 | usage guidelines..... | 20 |
| stateful firewall use with APPID..... | 677 | IPsec..... | 1488 |
| stateful NAT64 | | usage guidelines..... | 426, 429 |
| configuring..... | 89 | L2TP..... | 1488 |
| stateful-firewall-rule-sets statement | | usage guidelines..... | 644 |
| usage guidelines..... | 36 | NAT..... | 1489 |
| stateful-firewall-rules statement..... | 1486 | service sets..... | 1490 |
| usage guidelines..... | 36 | usage guidelines..... | 47 |
| statement | | stateful firewall..... | 1491 |
| flow monitoring | | usage guidelines..... | 334 |
| usage guidelines..... | 878, 1058 | usage guidelines..... | 516 |
| IPsec | | system log statement | |
| usage guidelines..... | 436 | NAT | |
| L2TP | | usage guidelines..... | 71 |
| usage guidelines..... | 652 | | |
| services | | T | |
| usage guidelines..... | 844 | t391 statement..... | 1615 |
| static destination address translation | | usage guidelines..... | 728 |
| configuring..... | 111, 183 | t392 statement..... | 1616 |
| statistics | | usage guidelines..... | 728 |
| AACL | | target statement..... | 1759 |
| clearing..... | 2083 | RPM..... | 1759 |
| active flow error..... | 2209 | tcp statement | |
| active flow instances..... | 2213 | RPM..... | 1759 |
| active flow memory utilization..... | 2223 | tcp-mss | |
| aggregated active flow..... | 2204 | statistics, displaying..... | 2031 |
| clearing inline flow instances..... | 2165 | tcp-mss statement..... | 1493 |
| | | tcp-tickles statement..... | 1286, 1777 |

- technical support
 - contacting JTAC.....lxi
- template-id statement.....1776
- template-refresh-rate statement.....1777
- templates
 - flow aggregation.....2208
- templates statement
 - video-monitoring.....1760
- term statement
 - AACL.....1588
 - usage guidelines.....661
 - CoS.....1494
 - HCM.....1498
 - IDS.....1495
 - usage guidelines.....355
 - IPsec.....1497
 - usage guidelines.....422
 - NAT.....1499
 - stateful firewall.....1500
 - usage guidelines.....331
 - usage guidelines.....514
- test services rpm rfc2544-benchmarking
 - command.....2293
- test statement
 - RPM.....1762
- test-interface (RFC 2544 Benchmarking)
 - RPM.....1764
- test-interval statement.....1765
- test-name (RFC 2544 Benchmarking).....1766
- tests (RFC 2544 Benchmarking).....1763
- then statement.....1501
 - AACL.....1589
 - usage guidelines.....661
 - HCM.....1501
 - IDS.....1502
 - usage guidelines.....355
 - IPsec.....1503
 - usage guidelines.....422
 - NAT.....1504
 - stateful firewall.....1505
 - usage guidelines.....331, 334
 - usage guidelines.....514
- threshold statement.....1506
 - usage guidelines.....358
- thresholds statement
 - RPM.....1767
- time-to-live threshold.....320
- timestamp option.....877, 1057
- topology
 - sample LFI and multilink bundle
 - network.....750, 755, 789
 - sample multilink bundle and LFI
 - network.....750, 755, 789
- trace-options
 - server (tracing flag).....50
 - timer-events (tracing flag).....50
- traceoptions statement
 - application identification.....1591
 - flow monitoring.....1768
 - IPsec.....1509
 - L-PDF.....1593
 - L2TP.....1511
 - RPM.....1769
 - security.....1507
 - services.....1515
- traceroute ALG
 - on the MS-MIC and MS-MPC.....322
- tracing flags
 - event policy
 - all.....50, 688
 - configuration.....50
 - database.....50
 - events.....50
 - policy.....50
 - server.....50
 - timer-events.....50
- tracing operations
 - adaptive services.....48
 - APPID.....686
 - RPM.....975, 1153
- traffic.....1253
 - inbound (decryption).....1257
 - IPsec, configuring.....1253
 - outbound (encryption).....1255
- traffic sampling
 - configuring.....871, 1051
 - disabling.....874, 1054, 1657
 - example configurations.....878, 1058
 - flow aggregation.....898, 1078
 - default values, option template
 - ID.....921, 1101
 - default values, template ID.....921, 1101
 - observation domain ID, version 9
 -918, 1098
 - option template ID, version 9 and
 - IPFIX.....921, 1101

| | |
|---|------------------|
| source ID, IPFIX..... | 918, 1098 |
| template ID, version 9 and IPFIX..... | 921, 1101 |
| FTP traffic..... | 880, 1060 |
| logging information output file..... | 878, 1058 |
| output files..... | 876, 1056 |
| SONET interfaces..... | 878, 1058 |
| traffic from single IP addresses..... | 879, 1059 |
| traffic-control-profiles statement | |
| usage guidelines..... | 532, 537 |
| transfer statement..... | 1770 |
| usage guidelines..... | 841 |
| transfer-log-archive statement..... | 1771 |
| usage guidelines..... | 842 |
| translated statement..... | 1517 |
| usage guidelines..... | 71 |
| translated-port statement | |
| NAT..... | 1518 |
| translation-type statement..... | 1519 |
| basic-nat-pt option..... | 143 |
| basic-nat44 option..... | 93 |
| basic-nat66 option..... | 99 |
| dnat-44 option, configuring..... | 111, 183 |
| dynamic-nat44, configuring..... | 191 |
| napt-44 option, configuring..... | 127 |
| napt-66 option, configuring..... | 131 |
| napt-pt option, example..... | 150 |
| stateful-nat64 option, configuring..... | 89 |
| usage guidelines..... | 71 |
| traps statement..... | 1772 |
| trigger-link-failure statement..... | 1517 |
| usage guidelines..... | 546 |
| trusted-ca statement..... | 1520 |
| usage guidelines..... | 433 |
| ttl statement | |
| DFC..... | 1773 |
| usage guidelines..... | 850 |
| tunnel..... | 1805 |
| ttl-threshold statement..... | 1521 |
| usage guidelines..... | 320 |
| tunnel interfaces | |
| configuration statements..... | 1213, 1219, 1239 |
| dynamic tunnels..... | 1245 |
| example configuration..... | 1218 |
| logical tunnels..... | 1221 |
| loopback tunnels..... | 1239 |
| multicast tunnels..... | 1219 |
| PIM tunnels..... | 1237 |
| unicast tunnels..... | 1213 |
| tunnel services interfaces..... | 2344 |
| tunnel statement..... | 1807 |
| encryption..... | 1806 |
| usage guidelines..... | 1251 |
| redundancy | |
| usage guidelines..... | 1260 |
| unicast | |
| usage guidelines..... | 1213 |
| tunnel-group statement..... | 1522 |
| usage guidelines..... | 641 |
| tunnel-mtu statement..... | 1523, 1524 |
| usage guidelines..... | 429, 436 |
| tunnel-timeout statement..... | 1525 |
| usage guidelines..... | 643 |
| tunnels | |
| definition..... | 1199 |
| GRE | |
| fragmentation of..... | 1216 |
| key number..... | 1215 |
| interface types..... | 1199 |
| IPv6-over-IPv4..... | 1211 |
| TWAMP server | |
| connections, clearing..... | 2168 |
| connections, displaying..... | 2281 |
| sessions, displaying..... | 2283 |
| twamp statement..... | 1774 |
| twamp-server statement..... | 1775 |
| twice NAT..... | 57 |
| twice-napt-44 option | |
| example..... | 187 |
| type statement..... | 1594 |
| APPID | |
| usage guidelines..... | 674 |
| type-of-service statement..... | 1595 |
| APPID | |
| usage guidelines..... | 674 |
| U | |
| udp statement | |
| RPM..... | 1778 |
| udp-tcp-port-swap (RFC 2544 | |
| Benchmarking)..... | 1686 |
| undirectional traffic support | |
| APPID..... | 685 |
| unicast tunnels..... | 1213 |
| unit statement | |
| aggregated Multiservices..... | 1527 |
| encryption..... | 1808 |
| usage guidelines..... | 1251 |

- flow monitoring.....1779
 - usage guidelines.....871, 1051
 - interfaces.....1528
 - link services.....1529, 1617
 - tunnel.....1809
 - usage guidelines.....1213
 - Universal Unique Identifier.....320
 - url statement.....1595
 - APPID
 - usage guidelines.....686
 - url-rule statement.....1526
 - url_identifier, statement.....1525, 1526
 - username statement
 - flow collection.....1780
 - usage guidelines.....842
 - uuid statement.....1530
 - usage guidelines.....320
- V**
- v6rd statement.....1531
 - usage guidelines.....245
 - var/log/sampled file.....878, 1058
 - var/tmp/sampled.pkts file.....876, 1056
 - variant statement.....1780
 - usage guidelines.....841
 - version statement
 - flow monitoring.....1781
 - IKE.....1532
 - usage guidelines.....411, 898, 1078
 - version-ipfix statement
 - usage guidelines.....890, 894, 1070, 1074
 - video monitoring
 - configuring.....1045, 1193
 - interface flow criteria.....1047, 1195
 - media delivery indexing.....1045, 1193
 - errors
 - clearing.....2169
 - displaying.....2285
 - flows
 - displaying.....2287
 - media delivery index See media delivery index
 - media delivery indexing
 - syslog messages.....1047, 1195
 - overview.....1043, 1191
 - platform support.....1043, 1191
 - statistics
 - clearing.....2170
 - displaying.....2291
 - video statement.....1532
 - usage guidelines.....517
 - video-monitoring statement
 - video-monitoring.....1782
 - virtual loopback tunnel
 - configuration guidelines.....1241
 - VRF table lookup
 - example configuration.....1241
 - virtual loopback tunnel interfaces
 - status information, displaying.....2340
 - voice services
 - bundles.....626
 - encapsulation.....625
 - example configuration.....627
 - interface type.....622
 - voice services interfaces
 - interleave fragments.....624
 - voice statement.....1533
 - usage guidelines.....517
- W**
- warm standby
 - AS PIC.....41
 - LSQ PIC.....548
 - warm-standby statement.....1533
 - warning (system logging severity level).....21, 47, 645
 - world-readable statement
 - flow monitoring.....1783
 - usage guidelines.....876, 1056
- Y**
- yellow-differential-delay statement.....1618
 - usage guidelines.....727

