



---

Junos<sup>®</sup> OS

# Adaptive Services Interfaces Feature Guide for Routing Devices

Release  
15.1



---

Modified: 2016-07-08

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://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.

*Junos<sup>®</sup> OS Adaptive Services Interfaces Feature Guide for Routing Devices*

15.1

Copyright © 2016, 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 . . . . .	xxxiii
	Documentation and Release Notes . . . . .	xxxiii
	Supported Platforms . . . . .	xxxiii
	Using the Examples in This Manual . . . . .	xxxiii
	Merging a Full Example . . . . .	xxxiv
	Merging a Snippet . . . . .	xxxiv
	Documentation Conventions . . . . .	xxxv
	Documentation Feedback . . . . .	xxxvii
	Requesting Technical Support . . . . .	xxxvii
	Self-Help Online Tools and Resources . . . . .	xxxvii
	Opening a Case with JTAC . . . . .	xxxviii
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Adaptive Services Overview . . . . .</b>	<b>3</b>
	Adaptive Services Overview . . . . .	3
	Packet Flow Through the Adaptive Services or Multiservices PIC . . . . .	5
<b>Chapter 2</b>	<b>Adaptive Services Configuration Overview . . . . .</b>	<b>7</b>
	Understanding Service Sets . . . . .	7
	Configuring Service Sets to be Applied to Services Interfaces . . . . .	9
	Configuring Interface Service Sets . . . . .	9
	Configuring Next-Hop Service Sets . . . . .	11
	Determining Traffic Direction . . . . .	12
	Interface Style Service Sets . . . . .	12
	Next-Hop Style Service Sets . . . . .	13
	Enabling the Reset of Service Sets for Aggregated Multiservices Interfaces . . . . .	13
	Configuring Service Rules . . . . .	14
	Configuring Service Set Limitations . . . . .	15
	Configuring Service Interface Pools . . . . .	16
	Enabling Services PICs to Accept Multicast Traffic . . . . .	17
	Applying Filters and Services to Interfaces . . . . .	17
	Configuring Service Filters . . . . .	18
	Example: Configuring Service Sets . . . . .	20
	Configuring AS or Multiservices PIC Redundancy . . . . .	20
	Examples: Configuring Services Interfaces . . . . .	23
	Configuring the Address and Domain for Services Interfaces . . . . .	24
	Configuring System Logging for Service Sets . . . . .	26
	Tracing Services PIC Operations . . . . .	27
	Configuring the Adaptive Services Log Filename . . . . .	28
	Configuring the Number and Size of Adaptive Services Log Files . . . . .	29

	Configuring Access to the Log File . . . . .	29
	Configuring a Regular Expression for Lines to Be Logged . . . . .	29
	Configuring the Trace Operations . . . . .	29
	Configuring Fragmentation Control for MS-DPC and MS-PIC Service Interfaces . . . . .	30
<b>Part 2</b>	<b>Translating IP Addresses Using NAT</b>	
<b>Chapter 3</b>	<b>NAT Overview . . . . .</b>	<b>35</b>
	Junos Address Aware Network Addressing Overview . . . . .	35
	NAT Concept and Facilities Overview . . . . .	36
	IPv4-to-IPv4 Basic NAT . . . . .	37
	Basic NAT . . . . .	37
	NAPT . . . . .	37
	Static Destination NAT . . . . .	37
	Twice NAT . . . . .	37
	IPv6 NAT . . . . .	38
	Application-Level Gateway (ALG) Support . . . . .	38
	NAT-PT with DNS ALG . . . . .	38
	Dynamic NAT . . . . .	39
	Stateful NAT64 . . . . .	39
	Dual-Stack Lite . . . . .	39
	Junos Address Aware Network Addressing Line Card Support . . . . .	40
	Junos OS Carrier-Grade NAT Implementation Overview . . . . .	40
	Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card . . . . .	41
	Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards . . . . .	43
	NAT Concept and Facilities Overview . . . . .	44
	IPv4-to-IPv4 Basic NAT . . . . .	45
	Basic NAT . . . . .	45
	NAPT . . . . .	45
	Static Destination NAT . . . . .	45
	Twice NAT . . . . .	46
	IPv6 NAT . . . . .	48
	Application-Level Gateway (ALG) Support . . . . .	48
	NAT-PT with DNS ALG . . . . .	48
	Dynamic NAT . . . . .	48
	Stateful NAT64 . . . . .	49
	Dual-Stack Lite . . . . .	49
<b>Chapter 4</b>	<b>NAT Configuration Overview . . . . .</b>	<b>51</b>
	Network Address Translation Configuration Overview . . . . .	51
	Configuring Source and Destination Addresses Network Address Translation Overview . . . . .	52
	Configuring Pools of Addresses and Ports for Network Address Translation Overview . . . . .	53
	Configuring NAT Pools . . . . .	53
	Preserve Range and Preserve Parity . . . . .	54
	Specifying Destination and Source Prefixes without Configuring a Pool . . . .	54



	Network Address Translation Rules Overview . . . . .	55
	Configuring Match Direction for NAT Rules . . . . .	56
	Configuring Match Conditions in NAT Rules . . . . .	56
	Configuring Actions in NAT Rules . . . . .	57
	Configuring Translation Types . . . . .	59
	Configuring Service Sets for Network Address Translation . . . . .	61
	Carrier-Grade NAT Implementation: Best Practices . . . . .	62
	Use APP and Round-Robin Address-Allocation . . . . .	63
	Do Not Use EIM with SIP . . . . .	63
	Do Not Use EIM with HTTP, DNS, or When Not Needed . . . . .	64
	Define PBA Blocks Based on User Profiles . . . . .	65
	Do Not Change the PBA Configuration on Running Systems . . . . .	65
	Do Not Allocate Excessively Large NAT Pools . . . . .	66
	Configure the System Log for PBA Only When Needed . . . . .	67
	Use Redundant Service PIC (RSP) Interfaces for Failover . . . . .	69
	Contain the Effects of Missing IP Fragments . . . . .	70
	Do Not Use Configurations Prone to Routing Loops . . . . .	70
<b>Chapter 5</b>	<b>Avoiding IPv4 Exhaustion Using Junos Address Aware Network Addressing and Stateful NAT64 . . . . .</b>	<b>73</b>
	Sample IPv6 Transition Scenarios . . . . .	73
	Example 1: IPv4 Depletion with a Non-IPv6 Access Network . . . . .	73
	Example 2: IPv4 Depletion with an IPv6 Access Network . . . . .	74
	Example 3: IPv4 Depletion for Mobile Networks . . . . .	75
	Configuring Stateful NAT64 . . . . .	75
<b>Chapter 6</b>	<b>Hiding Private Networks Using Static Source NAT . . . . .</b>	<b>79</b>
	Configuring Static Source Translation in IPv4 Networks . . . . .	79
	Configuring the NAT Pool and Rule . . . . .	79
	Configuring the Service Set for NAT . . . . .	81
	Configuring Trace Options . . . . .	83
	Sample Configuration - Static Source NAT Using a Static Pool With An Address Prefix And An Address Range . . . . .	84
	Sample Configuration - Static Source Nat for One-To-One Mapping Between a Private Subnet and a Public Subnet . . . . .	84
	Configuring Static Source Translation in IPv6 Networks . . . . .	85
	Configuring the NAT Pool and Rule . . . . .	86
	Configuring the Service Set for NAT . . . . .	87
	Configuring Trace Options . . . . .	88
	Example: Configuring Basic NAT44 . . . . .	89
	Example: Configuring NAT for Multicast Traffic . . . . .	91
	Rendezvous Point Configuration . . . . .	91
	Router 1 Configuration . . . . .	94
<b>Chapter 7</b>	<b>Making Private Servers Available Using Static Destination NAT . . . . .</b>	<b>97</b>
	Configuring Static Destination Address Translation in IPv4 Networks . . . . .	97

<b>Chapter 8</b>	<b>Allowing Components of a Private Network to Share a Single Address Using NAPT</b>	<b>103</b>
	Configuring Address Pools for Network Address Port Translation (NAPT)	
	Overview	103
	Round-Robin Allocation for NAPT	104
	Sequential Allocation for NAPT	104
	Preserve Parity and Preserve Range for NAPT	105
	Address Pooling and Endpoint Independent Mapping for NAPT	105
	Address Pooling	105
	Endpoint Independent Mapping and Endpoint Independent Filtering	106
	Port Block Allocation for NAPT	107
	Secured Port Block Allocation for NAPT	107
	Interim Logging for Port Block Allocation	108
	Deterministic Port Block Allocation for NAPT	108
	Understanding Deterministic Port Block Allocation Algorithms	108
	Deterministic Port Block Allocation Algorithm Usage	109
	Deterministic Port Block Allocation Restrictions	111
	Comparison of NAPT Implementation Methods	112
	Configuring Dynamic Source Address and Port Translation in IPv4 Networks	113
	Configuring Dynamic Source Address and Port Translation for IPv6 Networks	117
	Example: Configuring NAT with Port Translation	119
	Example: NAPT Configuration for the MS-MPC	120
	Example: Dynamic Source NAT as a Next-Hop Service	124
<b>Chapter 9</b>	<b>Securing Traffic Using NAT-PT and ALGs</b>	<b>127</b>
	ALGs Available by Default for Junos OS Address Aware NAT	127
	Configuring Protocol Translation Between IPv6 and IPv4 Networks -	
	NAT-PT	129
	Configuring the DNS ALG Application	130
	Configuring the NAT Pool and NAT Rule	130
	Configuring the Service Set for NAT	133
	Configuring Trace Options	134
	Example: Configuring NAT-PT	136
<b>Chapter 10</b>	<b>Reducing Traffic and Bandwidth Requirements Using Port Control Protocol</b>	<b>151</b>
	Port Control Protocol Overview	151
	Port Control Protocol Version 2	152
	Configuring Port Control Protocol	153
	Configuring PCP Server Options	153
	Configuring a PCP Rule	155
	Configuring a Service Set to Apply PCP	155
	SYSLOG Message Configuration	156
	Example: Configuring Port Control Protocol with NAPT44	156
<b>Chapter 11</b>	<b>Automatically Assigning Ports Using Port Block Allocation</b>	<b>163</b>
	Secured Port Block Allocation for NAPT	163
	Interim Logging for Port Block Allocation	163
	Guidelines for Configuring Interim Logging for Secured Port Block Allocation	164

	Secured Port Block Allocation for NAT44 and NAT64 Events on MS-MPCs and MS-MICs Overview . . . . .	167
	Guidelines for Configuring Secured Port Block Allocation for MS-MPCs and MS-MICs . . . . .	167
	Configuring Secured Port Block Allocation . . . . .	169
	Configuring Deterministic Port Block Allocation . . . . .	171
<b>Chapter 12</b>	<b>Connecting Specific Ports and Addresses Using Port Forwarding . . . . .</b>	<b>173</b>
	Configuring Port Forwarding for Static Destination Address Translation . . . . .	173
	Configuring Port Forwarding Without Destination Address Translation . . . . .	176
	Example: Configuring Port Forwarding with Twice NAT . . . . .	177
<b>Chapter 13</b>	<b>Allocating a Few Public Addresses to Many Private Hosts Using Dynamic NAT . . . . .</b>	<b>181</b>
	Configuring Dynamic Address-Only Source Translation in IPv4 Networks . . . . .	181
	Example: Dynamic Source NAT as a Next-Hop Service . . . . .	185
	Example: Assigning Addresses from a Dynamic Pool for Static Use . . . . .	187
<b>Chapter 14</b>	<b>Achieving Line-Rate, Low-Latency Translations Using Inline NAT . . . . .</b>	<b>189</b>
	Inline Network Address Translation Overview for MPC Types 1, 2, and 3 . . . . .	189
	Example: Configuring Inline Network Address Translation - Interface-Service Set . . . . .	191
<b>Chapter 15</b>	<b>Removing Address Dependency Using Network Prefix Translation for IPv6 Traffic . . . . .</b>	<b>199</b>
	Attaining Address Independence by Eliminating the Need for Advertising Internal Prefixes Using NPTv6 . . . . .	199
	Guidelines for Configuring Stateless Source Network Prefix Translation . . . . .	201
	Interoperation of Functionalities with Network Prefix Translation for IPv6 . . . . .	202
	Address Mapping Algorithm . . . . .	202
	Internal to External Translation . . . . .	203
	External to Internal Translation . . . . .	203
	Checksum-Neutral Translation . . . . .	203
	Multihoming . . . . .	203
	Hairpinning . . . . .	203
	Load Balancing . . . . .	204
	ICMPv6 for NPTv6 . . . . .	204
	Working of NPTv6 with Interface-Style and Next Hop-Style Service Sets . . . . .	204
	Example: Achieving Address Independence By Configuring Stateless Network Prefix Translation in IPv6 Networks by Using Interface-Style Service Sets . . . . .	205
	Example: Achieving Address Independence By Configuring Stateless Network Prefix Translation in IPv6 Networks by Using Next-Hop -Style Service Sets . . . . .	211
<b>Chapter 16</b>	<b>Monitoring NAT . . . . .</b>	<b>219</b>
	Configuring NAT Session Logs . . . . .	219
	Monitoring NAT Pool Usage . . . . .	220

<b>Part 3</b>	<b>Transitioning to IPv6 Using Softwires</b>	
<b>Chapter 17</b>	<b>Softwires Overview</b>	<b>223</b>
	Tunneling Services for IPv4-to-IPv6 Transition Overview	223
	6to4 Overview	223
	Basic 6to4	224
	6to4 Anycast	224
	6to4 Provider-Managed Tunnels	225
	DS-Lite Softwires—IPv4 over IPv6	225
	6rd Softwires—IPv6 over IPv4	226
<b>Chapter 18</b>	<b>Softwires Configuration Overview</b>	<b>229</b>
	Configuring Softwire Rules	229
	Configuring Service Sets for Softwire	230
<b>Chapter 19</b>	<b>Transitioning to IPv6 Using 6to4 Softwires</b>	<b>233</b>
	Configuring a 6to4 Provider-Managed Tunnel	233
<b>Chapter 20</b>	<b>Transitioning to IPv6 Using DS-Lite Softwires</b>	<b>237</b>
	Configuring a DS-Lite Softwire Concentrator	237
	Configuring IPv6 Multicast Interfaces	238
	Example: Basic DS-Lite Configuration	238
	Example: Configuring DS-Lite and 6rd in the Same Service Set	244
	Protecting CGN Devices Against Denial of Service (DOS) Attacks	251
	Mapping Refresh Behavior	251
	EIF Inbound Flow Limit	252
	DS-Lite Subnet Limitation	252
	DS-Lite Per Subnet Limitation Overview	252
	Configuring DS-Lite Per Subnet Session Limitation to Prevent Denial of Service Attacks	253
<b>Chapter 21</b>	<b>Transitioning to IPv6 Using 6rd Softwires</b>	<b>255</b>
	Configuring a 6rd Softwire Concentrator	255
	Configuring Stateful Firewall Rules for 6rd Softwire	256
	Example: Basic 6rd Configuration	257
	Inter-Chassis High Availability for MS-MIC and MS-MPC	262
	Inter-Chassis High Availability for Stateful Firewall and NAPT44 Overview (MS-MIC, MS-MPC)	262
	Configuring Inter-Chassis High Availability for Stateful Firewall and NAPT44 (MS-MPC, MS-MIC)	263
	Example: Inter-Chassis Stateful High Availability for NAT and Stateful Firewall (MS-MIC, MS-MPC)	264
	High Availability and Load Balancing for 6rd Softwires	274
	Load Balancing a 6rd Domain Across Multiple Services PICs	274
	Example: Load Balancing a 6rd Domain Across Multiple Services PICs	274
	Configuring High Availability for 6rd Using 6rd Anycast	279
	Configuring Inline 6rd	279
	Configuring the Bandwidth for Inline Services	280
	Configuring the Interfaces	280
	Configuring the Softwire Concentrator and Rule	282
	Configuring the Service Set	282

	Configuring the Routing Instance . . . . .	283
	Inline 6rd and 6to4 Configuration Guidelines . . . . .	283
	Examples: 6rd and 6to4 Configurations . . . . .	284
	Example: 6rd with Interface-Style Service Set Configuration . . . . .	284
	Example: 6rd with Next-Hop-Style Service Set Configuration . . . . .	285
	Example: 6rd Anycast Configuration . . . . .	287
	Example: Hairpinning Between 6rd Domains Configuration . . . . .	288
	Example: 6to4 Configuration . . . . .	290
<b>Chapter 22</b>	<b>Monitoring and Troubleshooting Softwires . . . . .</b>	<b>293</b>
	Ping and Traceroute for DS-Lite . . . . .	293
	Monitoring Softwire Statistics . . . . .	293
	Monitoring CGN, Stateful Firewall, and Softwire Flows . . . . .	295
<b>Part 4</b>	<b>Enabling Traffic to Pass Securely Using ALGs</b>	
<b>Chapter 23</b>	<b>ALG Overview . . . . .</b>	<b>299</b>
	ALG Descriptions . . . . .	299
	Supported ALGs . . . . .	299
	ALG Support Details . . . . .	300
	Basic TCP ALG . . . . .	301
	Basic UDP ALG . . . . .	302
	BOOTP . . . . .	302
	DCE RPC Services . . . . .	302
	DNS . . . . .	302
	FTP . . . . .	303
	H323 . . . . .	303
	ICMP . . . . .	304
	IIOP . . . . .	304
	IP . . . . .	304
	NetBIOS . . . . .	304
	NetShow . . . . .	305
	ONC RPC Services . . . . .	305
	PPTP . . . . .	305
	RealAudio . . . . .	305
	Sun RPC and RPC Portmap Services . . . . .	306
	RTSP . . . . .	307
	SIP . . . . .	308
	SNMP . . . . .	308
	SQLNet . . . . .	308
	TFTP . . . . .	309
	Traceroute . . . . .	309
	UNIX Remote-Shell Services . . . . .	309
	Winframe . . . . .	310
	Juniper Networks Defaults . . . . .	310
	Examples: Referencing the Preset Statement from the Junos OS Default Group . . . . .	320
	ALGs Available by Default for Junos OS Address Aware NAT . . . . .	321

<b>Chapter 24</b>	<b>ALGs Configuration Overview . . . . .</b>	<b>325</b>
	Configuring Application Sets . . . . .	325
	Configuring Application Protocol Properties . . . . .	325
	Configuring an Application Protocol . . . . .	326
	Configuring the Network Protocol . . . . .	328
	Configuring the ICMP Code and Type . . . . .	329
	Configuring Source and Destination Ports . . . . .	331
	Configuring the Inactivity Timeout Period . . . . .	334
	Configuring SIP . . . . .	334
	SIP ALG Interaction with Network Address Translation . . . . .	335
	Junos OS SIP ALG Limitations . . . . .	341
	Configuring an SNMP Command for Packet Matching . . . . .	342
	Configuring an RPC Program Number . . . . .	342
	Configuring the TTL Threshold . . . . .	342
	Configuring a Universal Unique Identifier . . . . .	342
	Examples: Configuring Application Protocols . . . . .	343
	Verifying the Output of ALG Sessions . . . . .	344
	FTP Example . . . . .	344
	Sample Output . . . . .	344
	FTP System Log Messages . . . . .	345
	Analysis . . . . .	345
	Troubleshooting Questions . . . . .	346
	RTSP ALG Example . . . . .	346
	Sample Output . . . . .	347
	Analysis . . . . .	347
	Troubleshooting Questions . . . . .	347
	System Log Messages . . . . .	349
	System Log Configuration . . . . .	349
	System Log Output . . . . .	349
	ICMP, Ping, and Traceroute ALGs for MS-MICs and MS-MPCs . . . . .	350
	Monitoring Port Control Protocol Operations . . . . .	351
<b>Part 5</b>	<b>Securing Content Using Junos Network Secure and IDS</b>	
<b>Chapter 25</b>	<b>Junos Network Secure Overview . . . . .</b>	<b>355</b>
	Junos Network Secure Overview . . . . .	355
	Stateful Firewall Support for Application Protocols . . . . .	356
	Stateful Firewall Anomaly Checking . . . . .	356
<b>Chapter 26</b>	<b>Junos Network Secure Configuration Overview . . . . .</b>	<b>359</b>
	Configuring Stateful Firewall Rules . . . . .	359
	Configuring Match Direction for Stateful Firewall Rules . . . . .	360
	Configuring Match Conditions in Stateful Firewall Rules . . . . .	360
	Configuring Actions in Stateful Firewall Rules . . . . .	362
	Configuring IP Option Handling . . . . .	362
	Configuring Stateful Firewall Rule Sets . . . . .	363
	Examples: Configuring Stateful Firewall Rules . . . . .	363
	Example: BOOTP and Broadcast Addresses . . . . .	367
	Example: Configuring Layer 3 Services and the Services SDK on Two PICs . . . . .	367

	Example: Virtual Routing and Forwarding (VRF) and Service Configuration . . .	379
<b>Chapter 27</b>	<b>IDS Configuration Overview . . . . .</b>	<b>383</b>
	Understanding SYN Cookie Protection . . . . .	383
	Configuring IDS Rules . . . . .	384
	Configuring Match Direction for IDS Rules . . . . .	385
	Configuring Match Conditions in IDS Rules . . . . .	386
	Configuring Actions in IDS Rules . . . . .	387
	Configuring IDS Rule Sets . . . . .	392
	Examples: Configuring IDS Rules . . . . .	392
<b>Chapter 28</b>	<b>Monitoring Junos Network Secure . . . . .</b>	<b>397</b>
	Monitoring Stateful Firewall Conversations . . . . .	397
	Monitoring CGN, Stateful Firewall, and Software Flows . . . . .	397
	Monitoring Global Stateful Firewall Statistics . . . . .	398
<b>Part 6</b>	<b>Creating Secure Tunnels Using Junos VPN Site Secure</b>	
<b>Chapter 29</b>	<b>Junos VPN Site Secure Overview . . . . .</b>	<b>401</b>
	Understanding Junos VPN Site Secure . . . . .	401
	IPsec . . . . .	401
	Security Associations . . . . .	402
	IKE . . . . .	402
	Comparison of IPsec on ES PICs and Junos VPN Site Secure on Multiservices Line Cards . . . . .	402
	Authentication Algorithms . . . . .	404
	Encryption Algorithms . . . . .	404
	IPsec Protocols . . . . .	406
	Supported IPsec and IKE Standards . . . . .	408
	IPsec Terms and Acronyms . . . . .	409
<b>Chapter 30</b>	<b>Junos VPN Site Secure Configuration Overview . . . . .</b>	<b>413</b>
	Minimum Security Association Configurations . . . . .	413
	Minimum Manual SA Configuration . . . . .	413
	Minimum Dynamic SA Configuration . . . . .	414
	Configuring Security Associations . . . . .	415
	Configuring Manual Security Associations . . . . .	415
	Configuring the Direction for IPsec Processing . . . . .	416
	Configuring the Protocol for a Manual IPsec SA . . . . .	417
	Configuring the Security Parameter Index . . . . .	417
	Configuring the Auxiliary Security Parameter Index . . . . .	418
	Configuring Authentication for a Manual IPsec SA . . . . .	418
	Configuring Encryption for a Manual IPsec SA . . . . .	419
	Configuring Dynamic Security Associations . . . . .	420
	Clearing Security Associations . . . . .	420
	Example: Configuring Manual SAs . . . . .	421
	Configuring IKE Proposals . . . . .	435
	Configuring the Authentication Algorithm for an IKE Proposal . . . . .	436
	Configuring the Authentication Method for an IKE Proposal . . . . .	436
	Configuring the Diffie-Hellman Group for an IKE Proposal . . . . .	437

Configuring the Encryption Algorithm for an IKE Proposal . . . . .	438
Configuring the Lifetime for an IKE SA . . . . .	438
Example: Configuring an IKE Proposal . . . . .	439
Configuring IKE Policies . . . . .	439
Configuring the IKE Phase . . . . .	441
Configuring the Mode for an IKE Policy . . . . .	441
Configuring the Proposals in an IKE Policy . . . . .	441
Configuring the Preshared Key for an IKE Policy . . . . .	441
Configuring the Local Certificate for an IKE Policy . . . . .	442
Configuring a Certificate Revocation List . . . . .	443
Configuring the Description for an IKE Policy . . . . .	443
Configuring Local and Remote IDs for IKE Phase 1 Negotiation . . . . .	443
Enabling Invalid SPI Recovery . . . . .	444
Example: Configuring an IKE Policy . . . . .	444
Configuring IPsec Proposals . . . . .	445
Configuring the Authentication Algorithm for an IPsec Proposal . . . . .	446
Configuring the Description for an IPsec Proposal . . . . .	448
Configuring the Encryption Algorithm for an IPsec Proposal . . . . .	448
Configuring the Lifetime for an IPsec SA . . . . .	448
Configuring the Protocol for a Dynamic SA . . . . .	449
Configuring IPsec Policies . . . . .	450
Configuring the Description for an IPsec Policy . . . . .	450
Configuring Perfect Forward Secrecy . . . . .	451
Configuring the Proposals in an IPsec Policy . . . . .	451
IPsec Policy for Dynamic Endpoints . . . . .	451
Example: Configuring an IPsec Policy . . . . .	452
Configuring IPsec Rules . . . . .	452
Configuring Match Direction for IPsec Rules . . . . .	454
Configuring Match Conditions in IPsec Rules . . . . .	454
Configuring Actions in IPsec Rules . . . . .	456
Enabling IPsec Packet Fragmentation . . . . .	457
Configuring Destination Addresses for Dead Peer Detection . . . . .	457
Configuring or Disabling IPsec Anti-Replay . . . . .	458
Enabling System Log Messages . . . . .	459
Specifying the MTU for IPsec Tunnels . . . . .	459
Configuring IPsec Rule Sets . . . . .	459
Service Sets . . . . .	460
Configuring IPsec Service Sets . . . . .	460
Configuring the Local Gateway Address for IPsec Service Sets . . . . .	461
IKE Addresses in VRF Instances . . . . .	462
Configuring IKE Access Profiles for IPsec Service Sets . . . . .	462
Configuring Certification Authorities for IPsec Service Sets . . . . .	463
Configuring or Disabling Antireplay Service . . . . .	463
Clearing the Don't-Fragment Bit . . . . .	464
Configuring Passive-Mode Tunneling . . . . .	465
Configuring the Tunnel MTU Value . . . . .	466
Tracing Junos VPN Site Secure Operations . . . . .	466
Disabling IPsec Tunnel Endpoint in Traceroute . . . . .	467
Tracing IPsec PKI Operations . . . . .	468



	Multitask Example: Configuring IPsec Services . . . . .	468
	Configuring the IKE Proposal . . . . .	469
	Configuring the IKE Policy (and Referencing the IKE Proposal) . . . . .	469
	Configuring the IPsec Proposal . . . . .	470
	Configuring the IPsec Policy (and Referencing the IPsec Proposal) . . . . .	471
	Configuring the IPsec Rule (and Referencing the IKE and IPsec Policies) . . . . .	471
	Configuring IPsec Trace Options . . . . .	473
	Configuring the Access Profile (and Referencing the IKE and IPsec Policies) . . . . .	473
	Configuring the Service Set (and Referencing the IKE Profile and the IPsec Rule) . . . . .	474
	Example: Configuring Junos VPN Site Secure on MS MIC and MS-MPC . . . . .	475
<b>Chapter 31</b>	<b>Enhancing Security with Static IPsec over VRF . . . . .</b>	<b>487</b>
	Example: Configuring Statically Assigned IPsec Tunnels over a VRF Instance . . . . .	487
<b>Chapter 32</b>	<b>Dynamically Assigning Tunnels Using Junos VPN Site Secure . . . . .</b>	<b>495</b>
	Configuring Dynamic Endpoints for IPsec Tunnels . . . . .	495
	Authentication Process . . . . .	496
	Implicit Dynamic Rules . . . . .	496
	Reverse Route Insertion . . . . .	497
	Configuring an IKE Access Profile . . . . .	497
	Referencing the IKE Access Profile in a Service Set . . . . .	499
	Configuring the Interface Identifier . . . . .	499
	Default IKE and IPsec Proposals . . . . .	500
	Requesting for and Installing a Digital Certificates on Your Router . . . . .	501
	Requesting a Digital Certificate—Manual Process . . . . .	501
	Example: Configuring Dynamically Assigned Policy Based Tunnels . . . . .	503
	Example: Configuring IKE Dynamic SAs . . . . .	509
	Example: IKE Dynamic SA Configuration with Digital Certificates . . . . .	525
<b>Chapter 33</b>	<b>Enabling IPsec for the Services SDK . . . . .</b>	<b>549</b>
	Configuring Junos VPN Site Secure or IPsec VPN . . . . .	549
<b>Part 7</b>	<b>Alleviating Congestion and Controlling Service Using CoS</b>	
<b>Chapter 34</b>	<b>Class of Service Overview . . . . .</b>	<b>553</b>
	Class of Service Overview . . . . .	553
<b>Chapter 35</b>	<b>Class of Service Configuration Overview . . . . .</b>	<b>555</b>
	Restrictions and Cautions for CoS Configuration on Services Interfaces . . . . .	555
	Configuring CoS Rules . . . . .	556
	Configuring Match Direction for CoS Rules . . . . .	557
	Configuring Match Conditions In CoS Rules . . . . .	557
	Configuring Actions in CoS Rules . . . . .	558
	Configuring Application Profiles for Use as CoS Rule Actions . . . . .	559
	Configuring Reflexive and Reverse CoS Rule Actions . . . . .	560
	Example: Configuring CoS Rules . . . . .	560
	Configuring CoS Rule Sets . . . . .	561
	Examples: Configuring CoS on Services Interfaces . . . . .	561

<b>Chapter 36</b>	<b>Configuring Class of Service on LSQ Interfaces . . . . .</b>	<b>565</b>
	Link Services Configuration for Junos Interfaces . . . . .	565
	Configuring CoS Scheduling Queues on Logical LSQ Interfaces . . . . .	566
	Configuring Scheduler Buffer Size . . . . .	568
	Configuring Scheduler Priority . . . . .	568
	Configuring Scheduler Shaping Rate . . . . .	568
	Configuring Drop Profiles . . . . .	569
	Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces . . . . .	570
	Configuring Link Services and CoS on Services PICs . . . . .	572
	Oversubscribing Interface Bandwidth on LSQ Interfaces . . . . .	575
	Examples: Oversubscribing an LSQ Interface . . . . .	578
	Configuring Guaranteed Minimum Rate on LSQ Interfaces . . . . .	580
	Example: Configuring Guaranteed Minimum Rate . . . . .	583
<b>Part 8</b>	<b>Configuring Interface Redundancy and Bundling on LSQ Interfaces</b>	
<b>Chapter 37</b>	<b>Overview . . . . .</b>	<b>587</b>
	Layer 2 Service Package Capabilities and Interfaces . . . . .	587
<b>Chapter 38</b>	<b>Configuring Interface Redundancy with SONET APS and Virtual Interfaces . . . . .</b>	<b>589</b>
	Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS . . . . .	589
	Configuring the Association between LSQ and SONET Interfaces . . . . .	590
	Configuring SONET APS Interoperability with Cisco Systems FRF.16 . . . . .	591
	Restrictions on APS Redundancy for LSQ Interfaces . . . . .	591
	Configuring LSQ Interface Redundancy in a Single Router Using SONET APS . . . . .	592
	Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces . . . . .	592
	Configuring Redundant Paired LSQ Interfaces . . . . .	593
	Restrictions on Redundant LSQ Interfaces . . . . .	594
	Configuring Link State Replication for Redundant Link PICs . . . . .	595
	Examples: Configuring Redundant LSQ Interfaces for Failure Recovery . . . . .	597
<b>Chapter 39</b>	<b>Enabling Bundling on LSQ Interfaces . . . . .</b>	<b>603</b>
	Inline MLPPP for WAN Interfaces Overview . . . . .	603
	Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces . . . . .	605
	Configuring Multiclass MLPPP on LSQ Interfaces . . . . .	606
	Enabling Inline LSQ Services . . . . .	607
	Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using MLPPP . . . . .	609
	Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP . . . . .	612
	Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.16 . . . . .	615
	Example: Configuring an LSQ Interface as an NxT1 Bundle Using FRF.16 . . . . .	618
	Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.15 . . . . .	620
	Configuring LSQ Interfaces for Single Fractional T1 or El Interfaces Using MLPPP and LFI . . . . .	621
	Example: Configuring an LSQ Interface for a Fractional T1 Interface Using MLPPP and LFI . . . . .	624

	Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12 .....	626
	Examples: Configuring an LSQ Interface for a Fractional T1 Interface Using FRF.12 .....	629
	Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP .....	633
	Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12 .....	635
	Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP .....	637
<b>Part 9</b>	<b>Enabling Load Balancing and High Availability Using Multiservices Interfaces</b>	
<b>Chapter 40</b>	<b>Enabling Load Balancing and High Availability Using Multiservices Interfaces .....</b>	<b>643</b>
	Understanding Aggregated Multiservices Interfaces .....	643
	Aggregated Multiservices Interface .....	643
	IPv6 Traffic on AMS Interfaces Overview .....	647
	Member Failure Options and High Availability Settings .....	648
	Configuring Load Balancing on AMS Infrastructure .....	649
	Configuring AMS Infrastructure .....	650
	Configuring High Availability .....	651
	Load Balancing Network Address Translation Flows .....	651
	Example: Configuring an Aggregated Multiservices Interface (AMS) .....	652
	Example: Configuring Next-Hop Style Services on an Aggregated Multiservices Interface .....	657
	Example: Configuring Static Source Translation on AMS Infrastructure .....	660
<b>Part 10</b>	<b>Handling VoIP, HTTP, and Layer 2 Traffic</b>	
<b>Chapter 41</b>	<b>Handling VoIP Traffic Using Voice Services .....</b>	<b>665</b>
	Voice Services Overview .....	665
	Configuring Services Interfaces for Voice Services .....	666
	Configuring the Logical Interface Address for the MLPPP Bundle .....	666
	Configuring Compression of Voice Traffic .....	667
	Configuring Delay-Sensitive Packet Interleaving .....	668
	Example: Configuring Compression of Voice Traffic .....	668
	Configuring Encapsulation for Voice Services .....	669
	Configuring Network Interfaces for Voice Services .....	670
	Configuring Voice Services Bundles with MLPPP Encapsulation .....	670
	Configuring the Compression Interface with PPP Encapsulation .....	670
	Examples: Configuring Voice Services .....	671
<b>Chapter 42</b>	<b>Tunneling PPP Packets Across a Network Using Layer 2 Tunneling .....</b>	<b>675</b>
	Layer 2 Tunneling Protocol Overview .....	675
	L2TP Services Configuration Overview .....	676
	L2TP Minimum Configuration .....	677
	Configuring L2TP Tunnel Groups .....	679
	Configuring Access Profiles for L2TP Tunnel Groups .....	680
	Configuring the Local Gateway Address and PIC .....	680
	Configuring Window Size for L2TP Tunnels .....	681

	Configuring Timers for L2TP Tunnels . . . . .	681
	Hiding Attribute-Value Pairs for L2TP Tunnels . . . . .	682
	Configuring System Logging of L2TP Tunnel Activity . . . . .	682
	Configuring the Identifier for Logical Interfaces that Provide L2TP Services . . . . .	684
	Example: Configuring Multilink PPP on a Shared Logical Interface . . . . .	684
	AS PIC Redundancy for L2TP Services . . . . .	686
	Examples: Configuring L2TP Services . . . . .	686
	Tracing L2TP Operations . . . . .	689
<b>Part 11</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 43</b>	<b>Configuration Statements . . . . .</b>	<b>695</b>
	[edit services application-identification] Hierarchy Level . . . . .	704
	IPsec Hierarchy Level . . . . .	706
	adaptive-services-pics . . . . .	709
	address (Interfaces) . . . . .	710
	address (Services NAT Pool) . . . . .	710
	address-allocation . . . . .	711
	address-range . . . . .	711
	aggregation . . . . .	712
	allow-ip-options . . . . .	713
	allow-multicast . . . . .	714
	allow-overlapping-nat-pools . . . . .	714
	anti-replay-window-size (Services IPsec VPN) . . . . .	714
	anti-replay-window-size (Services Service Set) . . . . .	715
	app-mapping-timeout . . . . .	716
	application . . . . .	717
	application-protocol . . . . .	718
	application-profile . . . . .	720
	application-set . . . . .	721
	application-sets (Services CoS) . . . . .	721
	application-sets (Services IDS) . . . . .	722
	application-sets (Services NAT) . . . . .	722
	application-sets (Services Stateful Firewall) . . . . .	723
	applications (Services ALGs) . . . . .	723
	applications (Services CoS) . . . . .	724
	applications (Services IDS) . . . . .	724
	applications (Services NAT) . . . . .	725
	applications (Services Stateful Firewall) . . . . .	725
	authentication . . . . .	726
	authentication-algorithm (Services IKE) . . . . .	727
	authentication-algorithm (Services IPsec) . . . . .	728
	authentication-method . . . . .	729
	auxiliary-spi . . . . .	730
	backup-remote-gateway . . . . .	730
	bundle . . . . .	731
	by-destination . . . . .	731
	by-pair . . . . .	732
	by-source . . . . .	733

bypass-traffic-on-exceeding-flow-limits	733
bypass-traffic-on-pic-failure	734
cgn-pic	734
cisco-interoperability	735
class	736
clear-dont-fragment-bit (Interfaces GRE Tunnels)	737
clear-dont-fragment-bit	737
clear-dont-fragment-bit (Services NAT Options)	738
clear-dont-fragment-bit (Services Service Set)	738
clear-ike-sas-on-pic-restart	739
clear-ipsec-sas-on-pic-restart	739
compression	740
compression-device (Interfaces)	740
copy-dont-fragment-bit (Services IPsec VPN)	741
copy-dont-fragment-bit (Services Set)	741
data (FTP)	742
dead-peer-detection (Services IPsec VPN)	742
description (Services IPsec VPN)	743
destination-address (Services CoS)	743
destination-address (Services IDS)	744
destination-address	744
destination-address (Services NAT)	745
destination-address (Services Stateful Firewall)	745
destination-address-range (Services IDS)	746
destination-address-range (Services NAT)	747
destination-address-range (Services Stateful Firewall)	748
destination-pool	748
destination-port	749
destination-port range	750
destination-prefix (Services IDS)	750
destination-prefix (Services NAT)	751
destination-prefix-ipv6	751
destination-prefix-list (Services CoS)	752
destination-prefix-list (Services IDS)	752
destination-prefix-list (Services NAT)	753
destination-prefix-list (Services Stateful Firewall)	753
destined-port	754
deterministic-port-block-allocation	755
dh-group	756
dial-options	757
direction	758
dns-alg-pool	758
dns-alg-prefix	759
drop-member-traffic (Aggregated Multiservices)	759
ds-lite	760
dscp	761
dynamic	761
ecmp-alb	762
ei-mapping-timeout	763

eif-flow-limit	763
enable-change-on-ams-redistribution	764
enable-rejoin (aggregated Multiservices)	765
encapsulation	766
encryption	767
encryption-algorithm	768
establish-tunnels	769
f-max-period	769
facility-override (Service Sets)	770
facility-override (System Log Reporting)	771
family (Aggregated Multiservices)	771
family (Interfaces)	772
family (Voice Services)	773
force-entry	774
forwarding-class (Services CoS)	774
forwarding-class (Services CoS Fragmentation Properties)	775
fragment-limit	775
fragment-threshold (Forwarding Class Maps)	776
fragment-threshold (Interfaces LSQ)	777
fragmentation-map	777
fragmentation-maps	778
from (Services CoS)	779
from (Services IDS)	780
from	781
from (Services HCM)	781
from (Services NAT)	782
from (Services Stateful Firewall)	783
ftp (Services CoS)	784
hash-keys (Aggregated Multiservices)	785
header-integrity-check	788
hello-interval	789
hide-avps	790
high-availability-options (aggregated Multiservices)	791
hint	792
host (L2TP)	792
host (service-set)	793
host (Services HCM)	794
hot-standby	794
icmp-code	795
icmp-type	795
ids-rules	796
ignore-entry	796
ike	797
ike-access-profile	798
inactivity-timeout	798
initiate-dead-peer-detection	799
input (Interfaces)	799
interface	800
interface-service	800

interfaces (Aggregated Multiservices) . . . . .	801
interfaces (Voice Services) . . . . .	802
interval . . . . .	802
ipsec . . . . .	803
ipsec-inside-interface . . . . .	803
ipsec-vpn-options . . . . .	804
ipsec-vpn-rules . . . . .	804
ipv6-multicast-interfaces . . . . .	805
l2tp-access-profile . . . . .	805
land-attack-check . . . . .	806
learn-sip-register . . . . .	806
lifetime-seconds . . . . .	807
link-layer-overhead . . . . .	807
load-balance . . . . .	808
load-balancing-options (Aggregated Multiservices) . . . . .	809
local-certificate . . . . .	810
local-gateway (IPSec) . . . . .	811
local-gateway (L2TP LNS) . . . . .	811
local-id . . . . .	812
log-prefix (L2TP) . . . . .	812
log-prefix (Services) . . . . .	813
logging (Services) . . . . .	813
logging (Services IDS) . . . . .	814
lsq-failure-options . . . . .	814
manual . . . . .	815
many-to-one (Aggregated Multiservices) . . . . .	816
mapping-refresh . . . . .	817
mapping-timeout . . . . .	818
match-direction (Services CoS) . . . . .	818
match-direction (Services IDS) . . . . .	819
match-direction . . . . .	819
match-direction (Services NAT) . . . . .	820
match-direction (Services Stateful Firewall) . . . . .	820
max-drop-flows . . . . .	821
max-flows . . . . .	822
max-sessions-per-subscriber . . . . .	823
maximum . . . . .	823
maximum-contexts . . . . .	824
maximum-send-window . . . . .	824
member-failure-options (Aggregated Multiservices) . . . . .	825
member-interface (Aggregated Multiservices) . . . . .	827
message-rate-limit . . . . .	828
mlfr-uni-nni-bundles-inline . . . . .	829
mode . . . . .	830
mss . . . . .	830
multi-link-layer-2-inline . . . . .	831
multilink-class . . . . .	831
multilink-max-classes . . . . .	832
nat-options . . . . .	832

nat-rules	833
next-hop-service	834
no-anti-replay	835
no-anti-replay (Services Service Set)	835
no-fragmentation	836
no-ipsec-tunnel-in-traceroute	836
no-per-unit-scheduler	837
no-termination-request	837
no-translation	838
output	838
overload-pool	839
overload-prefix	839
passive-mode-tunneling	840
pba-interim-logging-interval	841
per-unit-scheduler	842
perfect-forward-secrecy	843
pgcp	844
pgcp-rules	844
policy (Services IKE)	845
policy (IPsec)	846
pool	847
pool (Service Interface)	848
port (Services NAT)	849
port (Services Voice)	851
port (System Log Messages)	851
port-forwarding	852
port-forwarding-mappings	852
ports-per-session	853
post-service-filter	853
ppp-access-profile	854
pre-shared-key (Services IKE)	854
preserve-interface	855
primary (Adaptive Services Interfaces)	855
primary (Link Services IQ PIC Interfaces)	856
proposal (Services IKE)	856
proposal (Services IPsec VPN)	857
proposals	857
protocol (Applications)	858
protocol (IPSec)	859
ptsp-rules	859
queues	860
reassembly-timeout	860
receive-window	861
redistribute-all-traffic (Aggregated Multiservices)	861
redundancy-options (Adaptive Services Interfaces)	862
redundancy-options (Link Services IQ PIC Interfaces)	862
redundancy-options (MS-MIC, MS-MPC)	863
(reflexive   reverse)	864
rejoin-timeout (Aggregated Multiservices)	865



remote-gateway	865
remote-id	866
remotely-controlled	866
request-url	867
replicate-services (MS-MIC, MS-MPC)	868
respond-bad-spi (Services IKE Policy)	869
retransmit-interval (Services)	869
rpc-program-number	870
routing-engine-services	870
rtp	871
rule (Services CoS)	872
rule (Services IDS)	873
rule	875
rule (Services NAT)	877
rule (Services Stateful Firewall)	878
rule (Softwire)	879
rule-set (Services CoS)	879
rule-set (Services IDS)	880
rule-set	880
rule-set (Services NAT)	881
rule-set (Services Stateful Firewall)	881
rule-set (Softwire)	882
secondary (Adaptive Services Interfaces)	882
secondary (Link Services IQ PIC Interfaces)	883
secure-nat-mapping	883
secured-port-block-allocation	884
server (pcp)	886
service	887
service-domain	888
service-filter (Interfaces)	888
service-interface (Adaptive Services Interfaces)	889
service-interface (L2TP Processing)	889
service-interface-pools	890
service-set (Interfaces)	890
service-set (Services)	891
service-set-options	893
services (NAT)	893
session-limit	894
set-dont-fragment-bit (Services Set)	895
set-dont-fragment-bit (Services IPsec VPN)	895
sip-call-hold-timeout	896
sip	897
snmp-command	897
snmp-trap-thresholds	898
softwire-concentrator	899
softwire-options	900
softwire-rules	900
source-address (Service Sets)	901
source-address (Services CoS)	901

source-address (Services IDS) .....	902
source-address .....	902
source-address (Services NAT) .....	903
source-address (Services Stateful Firewall) .....	903
source-address-range (Services IDS) .....	904
source-address-range (Services NAT) .....	904
source-address-range (Services Stateful Firewall) .....	905
source-pool .....	905
source-port .....	906
source-prefix (Services IDS) .....	906
source-prefix (Services NAT) .....	907
source-prefix-ipv6 .....	907
source-prefix-list (Services CoS) .....	908
source-prefix-list (Services IDS) .....	908
source-prefix-list (Services NAT) .....	909
source-prefix-list (Services Stateful Firewall) .....	909
spi .....	910
stateful-firewall-rules .....	910
stateful-nat64 .....	911
syslog (Services CoS) .....	911
syslog (Services IDS) .....	912
syslog .....	912
syslog (Services L2TP) .....	913
syslog (Services NAT) .....	913
syslog (Services Service Set) .....	914
syslog (Services Stateful Firewall) .....	915
syn-cookie .....	916
tcp-mss .....	917
term (Services CoS) .....	918
term (Services IDS) .....	919
term .....	921
term (Services HCM) .....	922
term (Services NAT) .....	923
term (Services Stateful Firewall) .....	924
then (Services CoS) .....	925
then (Services HCM) .....	925
then (Services IDS) .....	926
then .....	927
then (Services NAT) .....	928
then (Services Stateful Firewall) .....	929
threshold (Services IPsec) .....	930
threshold (Services Logging and SYN-Cookie Defenses) .....	930
traceoptions (Security PKI) .....	931
traceoptions (Services IPsec VPN) .....	933
traceoptions (Services L2TP) .....	935
traceoptions (Services Logging) .....	939
translated .....	941
transport .....	941
trigger-link-failure .....	942

	translated-port . . . . .	942
	translation-type . . . . .	943
	trusted-ca . . . . .	945
	ttl-threshold . . . . .	945
	tunnel-group . . . . .	946
	tunnel-mtu (Services IPsec VPN) . . . . .	947
	tunnel-mtu (Services Service Set) . . . . .	948
	tunnel-timeout . . . . .	949
	url . . . . .	949
	url-list . . . . .	950
	url-rule . . . . .	950
	url-rule-set . . . . .	951
	unit (Aggregated Multiservices) . . . . .	951
	unit (Interfaces) . . . . .	952
	unit (Voice Services) . . . . .	953
	uuid . . . . .	954
	v6rd . . . . .	955
	version (IKE) . . . . .	956
	video . . . . .	956
	video (Application Profile) . . . . .	957
	voice . . . . .	957
	voice (Application Profile) . . . . .	958
	warm-standby . . . . .	958
<b>Chapter 44</b>	<b>Operational Commands . . . . .</b>	<b>959</b>
	clear services cos statistics . . . . .	963
	clear services crtp statistics . . . . .	964
	clear services ids . . . . .	965
	clear services ids destination-table . . . . .	966
	clear services ids pair-table . . . . .	967
	clear services ids source-table . . . . .	968
	clear services inline nat pool . . . . .	969
	clear services inline nat statistics . . . . .	970
	clear services inline software statistics . . . . .	971
	clear services ipsec-vpn certificates . . . . .	972
	clear services ipsec-vpn ike security-associations . . . . .	973
	clear services ipsec-vpn ipsec security-associations . . . . .	974
	clear services ipsec-vpn ipsec statistics . . . . .	975
	clear services l2tp destination . . . . .	976
	clear services l2tp destination statistics . . . . .	977
	clear services l2tp multilink . . . . .	978
	clear services l2tp session . . . . .	979
	clear services l2tp session statistics . . . . .	981
	clear services l2tp tunnel . . . . .	983
	clear services l2tp tunnel statistics . . . . .	985
	clear services nat flows . . . . .	987
	clear services nat mappings . . . . .	988
	clear services nat mappings app . . . . .	990
	clear services nat mappings eim . . . . .	991

clear services nat mappings pcp . . . . .	993
clear security pki ca-certificate . . . . .	995
clear security pki certificate-request . . . . .	996
clear security pki crl . . . . .	997
clear security pki key-pair . . . . .	998
clear security pki local-certificate . . . . .	999
clear services service-sets statistics integrity-drops . . . . .	1000
clear services service-sets statistics packet-drops . . . . .	1001
clear services service-sets statistics syslog . . . . .	1002
clear services sessions . . . . .	1003
clear services stateful-firewall flows . . . . .	1006
clear services stateful-firewall sip-call . . . . .	1008
clear services stateful-firewall sip-register . . . . .	1011
clear services stateful-firewall statistics . . . . .	1014
request interface (revert   switchover) (Adaptive Services) . . . . .	1015
request security pki ca-certificate enroll . . . . .	1016
request security pki ca-certificate load . . . . .	1017
request security pki ca-certificate verify . . . . .	1018
request security pki crl load . . . . .	1019
request security pki generate-certificate-request . . . . .	1020
request security pki generate-key-pair . . . . .	1022
request security pki local-certificate enroll . . . . .	1023
request security pki local-certificate generate-self-signed . . . . .	1025
request security pki local-certificate load . . . . .	1026
request security pki local-certificate verify . . . . .	1027
request services ipsec-vpn ipsec switch tunnel . . . . .	1028
show interfaces (Adaptive Services) . . . . .	1029
show interfaces (Link Services IQ) . . . . .	1037
show interfaces (Redundant Adaptive Services) . . . . .	1061
show interfaces (Redundant Link Services IQ) . . . . .	1063
show interfaces load-balancing . . . . .	1077
show interfaces redundancy . . . . .	1080
show security pki ca-certificate . . . . .	1082
show security pki certificate-request . . . . .	1086
show security pki crl . . . . .	1088
show security pki local-certificate . . . . .	1090
show services cos statistics . . . . .	1093
show services crtp . . . . .	1096
show services crtp flows . . . . .	1098
show services hcm statistics . . . . .	1100
show services ids . . . . .	1101
show services inline nat pool . . . . .	1109
show services inline nat statistics . . . . .	1111
show services inline software statistics . . . . .	1114
show services ipsec-vpn certificates . . . . .	1117
show services ipsec-vpn ike security-associations . . . . .	1120
show services ipsec-vpn ipsec security-associations . . . . .	1124
show services ipsec-vpn ipsec statistics . . . . .	1128
show services link-services cpu-usage . . . . .	1131

show services l2tp multilink . . . . .	1135
show services l2tp radius . . . . .	1139
show services l2tp session . . . . .	1143
show services l2tp summary . . . . .	1151
show services l2tp tunnel . . . . .	1156
show services l2tp user . . . . .	1162
show services nat deterministic-nat internal-host . . . . .	1166
show services nat deterministic-nat nat-port-block . . . . .	1168
show services nat ipv6-multicast-interfaces . . . . .	1169
show services nat mappings . . . . .	1171
show services nat pool . . . . .	1176
show services pcp statistics . . . . .	1181
show services service-sets cpu-usage . . . . .	1184
show services service-sets memory-usage . . . . .	1186
show services service-sets statistics integrity-drops . . . . .	1188
show services service-sets statistics packet-drops . . . . .	1192
show services service-sets statistics syslog . . . . .	1194
show services service-sets statistics tcp-mss . . . . .	1199
show services service-sets summary . . . . .	1200
show services sessions . . . . .	1202
show services software . . . . .	1210
show services software flows . . . . .	1211
show services software statistics . . . . .	1214
show services stateful-firewall conversations . . . . .	1220
show services stateful-firewall flow-analysis . . . . .	1224
show services stateful-firewall flows . . . . .	1228
show services stateful-firewall sip-call . . . . .	1234
show services stateful-firewall sip-register . . . . .	1239
show services stateful-firewall statistics . . . . .	1243
show services stateful-firewall statistics application-protocol sip . . . . .	1252
show services stateful-firewall subscriber-analysis . . . . .	1255

## Part 12

## Index

Index . . . . .	1261
-----------------	------



# List of Figures

<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Adaptive Services Overview</b>	<b>3</b>
	Figure 1: Packet Flow Through the Adaptive Services or MultiServices PIC	6
<b>Part 2</b>	<b>Translating IP Addresses Using NAT</b>	
<b>Chapter 3</b>	<b>NAT Overview</b>	<b>35</b>
	Figure 2: Dynamic NAT Flow	39
	Figure 3: Stateful NAT64 Flow	39
	Figure 4: DS-Lite Flow	40
	Figure 5: Dynamic NAT Flow	48
	Figure 6: Stateful NAT64 Flow	49
	Figure 7: DS-Lite Flow	49
<b>Chapter 5</b>	<b>Avoiding IPv4 Exhaustion Using Junos Address Aware Network Addressing and Stateful NAT64</b>	<b>73</b>
	Figure 8: IPv4 Depletion Solution - IPv4 Access Network	74
	Figure 9: IPv4 Depletion Solution - IPv6 Access Network	74
<b>Chapter 6</b>	<b>Hiding Private Networks Using Static Source NAT</b>	<b>79</b>
	Figure 10: Configuring NAT for Multicast Traffic	91
<b>Chapter 9</b>	<b>Securing Traffic Using NAT-PT and ALGs</b>	<b>127</b>
	Figure 11: Configuring DNS ALGs with NAT-PT Network Topology	137
<b>Chapter 10</b>	<b>Reducing Traffic and Bandwidth Requirements Using Port Control Protocol</b>	<b>151</b>
	Figure 12: Basic PCP NAPT44 Topology	152
	Figure 13: PCP with DS-Lite Plain Mode	152
	Figure 14: PCP with NAPT44	157
<b>Chapter 14</b>	<b>Achieving Line-Rate, Low-Latency Translations Using Inline NAT</b>	<b>189</b>
	Figure 15: Supported Inline NAT Types	190
	Figure 16: Deploy Inline NAT within L3VPN	192
<b>Part 3</b>	<b>Transitioning to IPv6 Using Softwires</b>	
<b>Chapter 17</b>	<b>Softwires Overview</b>	<b>223</b>
	Figure 17: 6rd Software Flow	226
<b>Chapter 20</b>	<b>Transitioning to IPv6 Using DS-Lite Softwires</b>	<b>237</b>
	Figure 18: DS-Lite Topology	239
<b>Chapter 21</b>	<b>Transitioning to IPv6 Using 6rd Softwires</b>	<b>255</b>

	Figure 19: Inter-Chassis High Availability Topology . . . . .	263
	Figure 20: Inter-Chassis High Availability Topology . . . . .	265
<b>Part 6</b>	<b>Creating Secure Tunnels Using Junos VPN Site Secure</b>	
<b>Chapter 29</b>	<b>Junos VPN Site Secure Overview . . . . .</b>	<b>401</b>
	Figure 21: AH Protocol . . . . .	406
	Figure 22: ESP Protocol . . . . .	407
<b>Chapter 30</b>	<b>Junos VPN Site Secure Configuration Overview . . . . .</b>	<b>413</b>
	Figure 23: Manual SA Topology . . . . .	422
<b>Chapter 32</b>	<b>Dynamically Assigning Tunnels Using Junos VPN Site Secure . . . . .</b>	<b>495</b>
	Figure 24: IPsec Dynamic Endpoint Tunneling Topology . . . . .	504
	Figure 25: IKE Dynamic SAs . . . . .	509
	Figure 26: MS PIC IKE Dynamic SA Topology Diagram . . . . .	526
<b>Part 8</b>	<b>Configuring Interface Redundancy and Bundling on LSQ Interfaces</b>	
<b>Chapter 39</b>	<b>Enabling Bundling on LSQ Interfaces . . . . .</b>	<b>603</b>
	Figure 27: Inline MLPPP for WAN Interfaces . . . . .	604



# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xxxiii</b>
	Table 1: Notice Icons . . . . .	xxxv
	Table 2: Text and Syntax Conventions . . . . .	xxxvi
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 2</b>	<b>Adaptive Services Configuration Overview</b> . . . . .	<b>7</b>
	Table 3: System Log Message Severity Levels . . . . .	26
	Table 4: Adaptive Services Tracing Flags . . . . .	30
<b>Part 2</b>	<b>Translating IP Addresses Using NAT</b>	
<b>Chapter 3</b>	<b>NAT Overview</b> . . . . .	<b>35</b>
	Table 5: Carrier-Grade NAT—Feature Comparison by Platform . . . . .	41
	Table 6: Carrier-Grade NAT Translation Types . . . . .	43
<b>Chapter 8</b>	<b>Allowing Components of a Private Network to Share a Single Address Using NAPT</b> . . . . .	<b>103</b>
	Table 7: Deterministic Port Block Allocation Commit Constraints . . . . .	112
	Table 8: Comparison of NAPT Implementation Methods . . . . .	112
<b>Chapter 9</b>	<b>Securing Traffic Using NAT-PT and ALGs</b> . . . . .	<b>127</b>
	Table 9: ALGs Available by Default . . . . .	127
<b>Part 4</b>	<b>Enabling Traffic to Pass Securely Using ALGs</b>	
<b>Chapter 23</b>	<b>ALG Overview</b> . . . . .	<b>299</b>
	Table 10: ALGs Supported by Junos OS . . . . .	299
	Table 11: RealAudio Product Port Usage . . . . .	305
	Table 12: Supported RPC Services . . . . .	306
	Table 13: ALGs Available by Default . . . . .	322
<b>Chapter 24</b>	<b>ALGs Configuration Overview</b> . . . . .	<b>325</b>
	Table 14: Application Protocols Supported by Services Interfaces . . . . .	326
	Table 15: Network Protocols Supported by Services Interfaces . . . . .	328
	Table 16: ICMP Codes and Types Supported by Services Interfaces . . . . .	330
	Table 17: Port Names Supported by Services Interfaces . . . . .	331
	Table 18: Requesting Messages with NAT Table . . . . .	339
<b>Part 5</b>	<b>Securing Content Using Junos Network Secure and IDS</b>	
<b>Chapter 26</b>	<b>Junos Network Secure Configuration Overview</b> . . . . .	<b>359</b>
	Table 19: IP Option Values . . . . .	363

<b>Part 6</b>	<b>Creating Secure Tunnels Using Junos VPN Site Secure</b>	
<b>Chapter 29</b>	<b>Junos VPN Site Secure Overview</b>	<b>401</b>
	Table 20: Statement Equivalents for ES and AS Interfaces	403
<b>Chapter 32</b>	<b>Dynamically Assigning Tunnels Using Junos VPN Site Secure</b>	<b>495</b>
	Table 21: Default IKE and IPsec Proposals for Dynamic Negotiations	500
<b>Part 9</b>	<b>Enabling Load Balancing and High Availability Using Multiservices Interfaces</b>	
<b>Chapter 40</b>	<b>Enabling Load Balancing and High Availability Using Multiservices Interfaces</b>	<b>643</b>
	Table 22: Key Configuration Statements Used in this Example	655
<b>Part 10</b>	<b>Handling VoIP, HTTP, and Layer 2 Traffic</b>	
<b>Chapter 42</b>	<b>Tunneling PPP Packets Across a Network Using Layer 2 Tunneling</b>	<b>675</b>
	Table 23: System Log Message Severity Levels	682
<b>Part 11</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 43</b>	<b>Configuration Statements</b>	<b>695</b>
	Table 24: Hash Keys Supported for AMS for Service Applications	786
	Table 25: Behavior of Member Interface After One Multiservices PIC Fails	825
	Table 26: Behavior of Member Interface After Two Multiservices PICs Fail	826
<b>Chapter 44</b>	<b>Operational Commands</b>	<b>959</b>
	Table 27: clear services nat flows Output Fields	987
	Table 28: clear services nat mappings Output Fields	988
	Table 29: clear services nat mappings app Output Fields	990
	Table 30: clear services nat mappings eim Output Fields	991
	Table 31: clear services nat mappings pcp Output Fields	993
	Table 32: clear services sessions Output Fields	1005
	Table 33: clear services stateful-firewall flows Output Fields	1007
	Table 34: clear services stateful-firewall sip-call Output Fields	1010
	Table 35: clear services stateful-firewall sip-register Output Fields	1013
	Table 36: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields	1029
	Table 37: show interfaces (Link Services IQ) Output Fields	1038
	Table 38: show interfaces (Redundant Link Services IQ) Output Fields	1063
	Table 39: Aggregated Multiservices show interfaces load-balancing Output Fields	1077
	Table 40: show interfaces redundancy Output Fields	1080
	Table 41: show security pki ca-certificate Output Fields	1082
	Table 42: show security pki certificate-request Output Fields	1086
	Table 43: show security pki crl Output Fields	1088
	Table 44: show security pki local-certificate Output Fields	1090
	Table 45: show services cos statistics Output Fields	1093
	Table 46: show services crtp Output Fields	1096
	Table 47: show services crtp flows Output Fields	1098

Table 48: show services hcm statistics rule Output Fields . . . . .	1100
Table 49: show services ids Output Fields . . . . .	1102
Table 50: show services inline nat pool Output Fields . . . . .	1109
Table 51: show services inline nat statistics Output Fields . . . . .	1111
Table 52: show services inline software statistics Output Fields . . . . .	1114
Table 53: show services ipsec-vpn certificates Output Fields . . . . .	1117
Table 54: show services ipsec-vpn ike security-associations Output Fields . . . .	1120
Table 55: show services ipsec-vpn ipsec security-associations Output Fields . .	1124
Table 56: show services ipsec-vpn ipsec statistics Output Fields . . . . .	1128
Table 57: show services link-services cpu-usage Output Fields . . . . .	1131
Table 58: show services l2tp multilink Output Fields . . . . .	1135
Table 59: show services l2tp radius Output Fields . . . . .	1139
Table 60: show services l2tp session Output Fields . . . . .	1144
Table 61: show services l2tp summary Output Fields . . . . .	1151
Table 62: show services l2tp tunnel Output Fields . . . . .	1157
Table 63: show services l2tp user Output Fields . . . . .	1162
Table 64: show services nat deterministic-nat internal-host Output Fields . .	1166
Table 65: show services nat deterministic-nat nat-port-block Output Fields . .	1168
Table 66: show services nat ipv6-multicast-interfaces Output Fields . . . . .	1169
Table 67: show services nat mappings Output Fields . . . . .	1172
Table 68: show services nat pool Output Fields . . . . .	1177
Table 69: show services pcg statistics Output Fields . . . . .	1181
Table 70: show services service-sets cpu-usage Output Fields . . . . .	1184
Table 71: show services service-sets memory-usage Output Fields . . . . .	1186
Table 72: show services service-sets integrity-drops Output Fields . . . . .	1188
Table 73: show services service-sets packet-drops Output Fields . . . . .	1192
Table 74: show services service-sets statistics syslog Output Fields . . . . .	1194
Table 75: show services service-sets statistics tcp-mss Output Fields . . . . .	1199
Table 76: show services service-sets summary Output Fields . . . . .	1200
Table 77: show services sessions Output Fields . . . . .	1205
Table 78: show-services-software Output Fields . . . . .	1210
Table 79: show services software flows Output Fields . . . . .	1212
Table 80: command-name Output Fields . . . . .	1214
Table 81: show services stateful-firewall conversations Output Fields . . . . .	1222
Table 82: show services stateful-firewall flow-analysis Output Fields . . . . .	1224
Table 83: show services stateful-firewall flows Output Fields . . . . .	1230
Table 84: show services stateful-firewall sip-call Output Fields . . . . .	1236
Table 85: show services stateful-firewall sip-register Output Fields . . . . .	1241
Table 86: show services stateful-firewall statistics Output Fields . . . . .	1243
Table 87: show services stateful-firewall statistics application-protocol-sip Output Fields . . . . .	1252
Table 88: show services stateful-firewall subscriber-analysis Output Fields . .	1255



# About the Documentation

- Documentation and Release Notes on page xxxiii
- Supported Platforms on page xxxiii
- Using the Examples in This Manual on page xxxiii
- Documentation Conventions on page xxxv
- Documentation Feedback on page xxxvii
- Requesting Technical Support on page xxxvii

## Documentation and Release Notes

---

To obtain the most current version of all Juniper Networks<sup>®</sup> 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 xxxv 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 xxxvi defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces or emphasizes important new terms.</li> <li>Identifies guide names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS CLI User Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>
<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@# <b>set system domain-name</b> <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"> <li>To configure a stub area, include the <b>stub</b> statement at the [edit protocols ospf area area-id] hierarchy level.</li> <li>The console port is labeled <b>CONSOLE</b>.</li> </ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub &lt;default-metric metric&gt;;</b>
(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.	<b>broadcast   multicast</b>  <b>(string1   string2   string3)</b>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members [ community-ids ]</b>
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
<b>Bold text like this</b>	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> <li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li> <li>To cancel the configuration, click <b>Cancel</b>.</li> </ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## 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 of the Juniper Networks TechLibrary 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 <http://www.juniper.net/techpubs/feedback/>.
- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto: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 Partner Support Service 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

# Overview

- [Adaptive Services Overview on page 3](#)
- [Adaptive Services Configuration Overview on page 7](#)



## CHAPTER 1

# Adaptive Services Overview

- [Adaptive Services Overview on page 3](#)
- [Packet Flow Through the Adaptive Services or Multiservices PIC on page 5](#)

## 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*.

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 649](#) 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*
  - [Packet Flow Through the Adaptive Services or Multiservices PIC on page 5](#)
  - *Enabling Service Packages*
  - *Services Configuration Procedure*
  - *Supported Platforms*

## 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 6](#). (You can configure a service set as either an interface service set or a next-hop service set.)

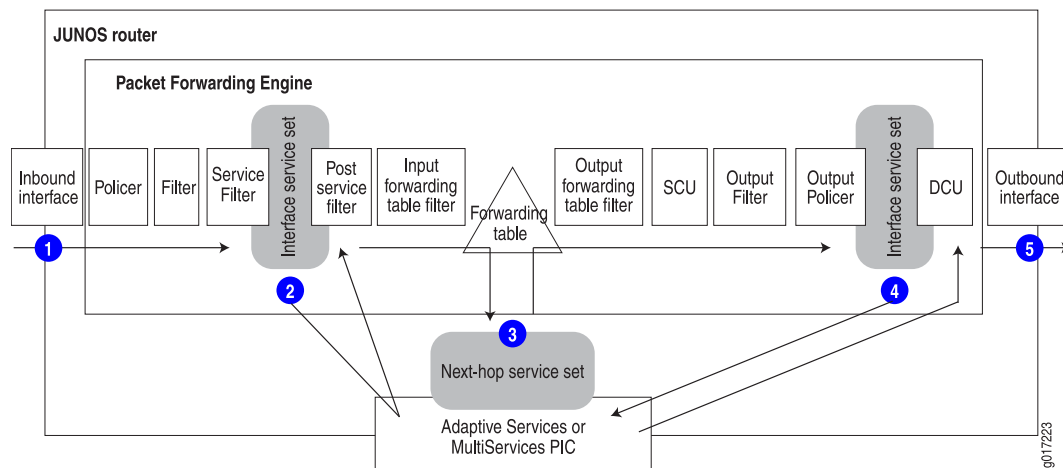
1. Packets enter the router on the inbound interface.
2. A policer, filter, service filter, service set, postservice filter, and input forwarding-table filter are applied sequentially to the traffic; these are all optional items in the configuration. If an interface service set is applied, the packets 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*
- [Adaptive Services Overview on page 3](#)
- *Supported Platforms*
- *Services Configuration Procedure*



## CHAPTER 2

# Adaptive Services Configuration Overview

- [Understanding Service Sets on page 7](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 9](#)
- [Enabling the Reset of Service Sets for Aggregated Multiservices Interfaces on page 13](#)
- [Configuring Service Rules on page 14](#)
- [Configuring Service Set Limitations on page 15](#)
- [Configuring Service Interface Pools on page 16](#)
- [Enabling Services PICs to Accept Multicast Traffic on page 17](#)
- [Applying Filters and Services to Interfaces on page 17](#)
- [Example: Configuring Service Sets on page 20](#)
- [Configuring AS or Multiservices PIC Redundancy on page 20](#)
- [Examples: Configuring Services Interfaces on page 23](#)
- [Configuring the Address and Domain for Services Interfaces on page 24](#)
- [Configuring System Logging for Service Sets on page 26](#)
- [Tracing Services PIC Operations on page 27](#)
- [Configuring Fragmentation Control for MS-DPC and MS-PIC Service Interfaces on page 30](#)

## 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 9](#)
- [Configuring Service Rules on page 14](#)
- [Configuring IPsec Service Sets on page 460](#)
- [Configuring Service Set Limitations on page 15](#)
- [Configuring System Logging for Service Sets on page 26](#)
- [Enabling Services PICs to Accept Multicast Traffic on page 17](#)
- [Tracing Services PIC Operations on page 27](#)
- [Example: Configuring Service Sets on page 20](#)

## 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 9](#)
- [Configuring Next-Hop Service Sets on page 11](#)
- [Determining Traffic Direction on page 12](#)

### 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 20](#).



**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 7](#)
- [Configuring Service Rules on page 14](#)
- [Configuring IPsec Service Sets on page 460](#)
- [Configuring Service Set Limitations on page 15](#)
- [Configuring System Logging for Service Sets on page 26](#)
- [Example: Configuring Service Sets on page 20](#)

## Enabling the Reset of Service Sets for Aggregated Multiservices Interfaces

Starting in Junos OS Release 15.1, for service sets associated with aggregated multiservices (AMS) interfaces, you can configure the **enable-change-on-ams-redistribution** statement at the **[edit services service-set service-set-name service-set-options]** hierarchy level to

enable the service set to be bounced (reset) for dynamic NAT scenarios (dynamic NAT, NAT64, and NAT44) when a member interface of an AMS bundle rejoins or a member interface failure occurs. When a member interface fails, the application resources (NAT pool in the case of dynamic NAT scenarios) and traffic load need to be rebalanced. For application resources to be rebalanced, which is the NAT pool for dynamic NAT environments, the NAT pool is split and allocated by the service PIC daemon (spd).

In real-world network conditions, the splitting of a NAT pool is catastrophic for the sessions that are already established. Therefore, if the count of active members of an AMS bundle changes, the service-set is bounced for resplitting the pool. The bouncing of the service set causes all service-set sessions to be reset, which results in an expensive operation when the number of service-sets is high. By default, the bouncing of service sets and splitting of a NAT pool is disabled in dynamic NAT conditions, such as the translation type being `dynamic-nat44`, `napt-44`, and `nat64`, which enables backward compatibility with earlier Junos OS releases.

To prevent the bouncing of a service set for resplitting the NAT pool and disabling the splitting of the NAT pool in dynamic NAT environments, you need not include the **enable-change-on-ams-redistribution** statement at the **[edit services service-set service-set-name service-set-options]** hierarchy level. When you enable this functionality by including this statement, the NAT pool is resplit (pool address spaces might be wasted) and the service set is bounced, which causes a minimal disruption to traffic. Traffic disruption is limited to sessions that are assigned to the member interfaces of an AMS bundle that failed or rejoined.

To enable the redistribution of application resources and traffic load, include the **enable-change-on-ams-redistribution** statement at the **[edit services service-set service-set-name service-set-options]** hierarchy level.

```
[edit services service-set service-set-name service-set-options]
user@host# set enable-change-on-ams-redistribution
```

**Related  
Documentation**

- [enable-change-on-ams-redistribution on page 764](#)

---

## 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 384](#).
- 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 software rules for DS-Lite or 6rd softwires at the **[edit services software]** hierarchy level; for more information, see “[Configuring Software Rules](#)” on page 229.
- 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);
([ software-rules rule-names ] | software-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 7](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 9](#)
- [Configuring Service Set Limitations on page 15](#)
- [Configuring System Logging for Service Sets on page 26](#)

## 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 392](#).



**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 7](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 9](#)
- [Configuring Service Rules on page 14](#)
- [Configuring System Logging for Service Sets on page 26](#)

---

## Configuring Service Interface Pools

To configure a service interface pool, include the following statements at the **[edit services service-interface-pools]** hierarchy level:

```
[edit services service-interface-pools]
pool pool-name {
  interface interface-name.unit-number;
}
```

**Related Documentation**

- [Configuring Service Sets to be Applied to Services Interfaces on page 9](#)

## 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 7](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 9](#)
- [Configuring Service Rules on page 14](#)
- [Example: Configuring Service Sets on page 20](#)
- [Example: Configuring NAT for Multicast Traffic on page 91](#)

## 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 9](#).



**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 23.

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](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 9](#)
- [Examples: Configuring Services Interfaces on page 23](#)

## 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 7](#)
  - [Configuring Service Sets to be Applied to Services Interfaces on page 9](#)

## 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 592.

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*.



**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 [“Configuring Security Associations” on page 415](#).

- 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*
  - [Examples: Configuring Services Interfaces on page 23](#)
  - [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 652](#)

## 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](#)
- [Configuring the Address and Domain for Services Interfaces on page 24](#)
- [Configuring Default Timeout Settings for Services Interfaces](#)
- [Configuring System Logging for Services Interfaces](#)
- [Applying Filters and Services to Interfaces on page 17](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 652](#)

---

## 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 9.

#### Related Documentation

- [Configuring Default Timeout Settings for Services Interfaces](#)
- [Configuring System Logging for Services Interfaces](#)
- [Examples: Configuring Services Interfaces on page 23](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 652](#)

## 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.



**NOTE:** Junos OS does not support the exporting of system log messages to an external system log server through the fxp.0 interface; this is because the high transmission rate of system log messages and the limited bandwidth of the fxp.0 interface can cause several problems. The external system log server must be reachable through a routable interface.

Table 3 on page 26 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 3: System Log Message Severity Levels**

Severity Level	Description
<b>any</b>	Includes all severity levels
<b>emergency</b>	System panic or other condition that causes the router to stop functioning
<b>alert</b>	Conditions that require immediate correction, such as a corrupted system database
<b>critical</b>	Critical conditions, such as hard drive errors

Table 3: System Log Message Severity Levels (*continued*)

Severity Level	Description
<b>error</b>	Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels
<b>warning</b>	Conditions that warrant monitoring
<b>notice</b>	Conditions that are not errors but might warrant special handling
<b>info</b>	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 [System Log Explorer](#).

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 7](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 9](#)
- [Tracing Services PIC Operations on page 27](#)

## 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.2**, 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 [System Log Explorer](#).)
- 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 28](#)
- [Configuring the Number and Size of Adaptive Services Log Files on page 29](#)
- [Configuring Access to the Log File on page 29](#)
- [Configuring a Regular Expression for Lines to Be Logged on page 29](#)
- [Configuring the Trace Operations on page 29](#)

## 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 4 on page 30 describes the meaning of the adaptive services tracing flags.

**Table 4: Adaptive Services Tracing Flags**

Flag	Description	Default Setting
<b>all</b>	Trace all operations.	Off
<b>command-queued</b>	Trace command enqueue events.	Off
<b>config</b>	Log reading of the configuration at the <b>[edit services]</b> hierarchy level.	Off
<b>handshake</b>	Trace handshake events.	Off
<b>init</b>	Trace initialization events.	Off
<b>interfaces</b>	Trace interface events.	Off
<b>mib</b>	Trace GGSN SNMP MIB events.	Off
<b>removed-client</b>	Trace client cleanup events.	Off
<b>show</b>	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 7](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 9](#)
- [Configuring System Logging for Service Sets on page 26](#)

## Configuring Fragmentation Control for MS-DPC and MS-PIC Service Interfaces

Two configuration options are available to prevent excessive consumption of computational CPU cycles on a services PIC caused by the handling of large numbers of fragmented packets. Such fragment handling can be exploited in DOS attacks. The **fragment-limit** option establishes a maximum number of fragments for a packet. When this number is exceeded, the packet is dropped. The **reassembly-timeout** specifies the



maximum time from the receipt of the first and latest fragments in a packet. When the number is exceeded, the packet is dropped.

To configure fragmentation control for MS-DPC and MS-PIC service interfaces:

1. In configuration mode, go to the **[edit interfaces *interface-name* services-options** hierarchy level.

**edit interfaces *interface-name* services-options**

2. Configure the fragment limit.

**[ edit services *interface-name* services-options]  
set fragment-limit *number-of-fragments***

3. Configure the reassembly timeout.

**[ edit services *interface-name* services-options]  
set reassembly-timeout *number-of-fragments***



## PART 2

# Translating IP Addresses Using NAT

- [NAT Overview on page 35](#)
- [NAT Configuration Overview on page 51](#)
- [Avoiding IPv4 Exhaustion Using Junos Address Aware Network Addressing and Stateful NAT64 on page 73](#)
- [Hiding Private Networks Using Static Source NAT on page 79](#)
- [Making Private Servers Available Using Static Destination NAT on page 97](#)
- [Allowing Components of a Private Network to Share a Single Address Using NAT on page 103](#)
- [Securing Traffic Using NAT-PT and ALGs on page 127](#)
- [Reducing Traffic and Bandwidth Requirements Using Port Control Protocol on page 151](#)
- [Automatically Assigning Ports Using Port Block Allocation on page 163](#)
- [Connecting Specific Ports and Addresses Using Port Forwarding on page 173](#)
- [Allocating a Few Public Addresses to Many Private Hosts Using Dynamic NAT on page 181](#)
- [Achieving Line-Rate, Low-Latency Translations Using Inline NAT on page 189](#)
- [Removing Address Dependency Using Network Prefix Translation for IPv6 Traffic on page 199](#)
- [Monitoring NAT on page 219](#)



## CHAPTER 3

# NAT Overview

- [Junos Address Aware Network Addressing Overview on page 35](#)
- [Junos OS Carrier-Grade NAT Implementation Overview on page 40](#)
- [Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 41](#)
- [Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards on page 43](#)

### 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 36](#)
- [IPv4-to-IPv4 Basic NAT on page 37](#)
- [Static Destination NAT on page 37](#)
- [Twice NAT on page 37](#)
- [IPv6 NAT on page 38](#)
- [Application-Level Gateway \(ALG\) Support on page 38](#)
- [NAT-PT with DNS ALG on page 38](#)
- [Dynamic NAT on page 39](#)
- [Stateful NAT64 on page 39](#)

- [Dual-Stack Lite on page 39](#)
- [Junos Address Aware Network Addressing Line Card Support on page 40](#)

## 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 37](#).
- 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 39](#).
  - 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 37](#).
- 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 37](#).
- 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 129](#), [“NAT-PT with DNS ALG” on page 38](#), and [“Stateful NAT64” on page 39](#).
- 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 223](#).

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 41](#), 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 55](#)

## 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 39](#).

**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 39](#).

**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 40](#).

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 41](#)
- [ALGs Available by Default for Junos OS Address Aware NAT on page 127](#)

## 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 [“Junos Address Aware Network Addressing Overview” on page 35](#).
- 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 [“Junos Address Aware Network Addressing Overview” on page 35](#).

- 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 189](#).

**Related Documentation**

- [Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 41](#)
- [Carrier-Grade NAT Implementation: Best Practices on page 62](#)
- [Example: Configuring Basic NAT44 on page 89](#)

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

[Table 5 on page 41](#) summarizes feature differences among the Junos OS carrier-grade NAT implementations.

**Table 5: 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	yes	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 189](#)

Table 5: 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	yes	yes  <small>NOTE: Twice NAT can be implemented indirectly. See "Inline Network Address Translation Overview for MPC Types 1, 2, and 3" on page 189</small>
NAPT - Preserve Parity and Range	yes	yes	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"> <li>• FTP</li> <li>• TFTP</li> <li>• SIP</li> <li>• RTSP</li> <li>• PPPT</li> </ul>			
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
No translation	yes	yes	yes

Table 6 on page 43 summarizes availability of translation types by type of line card.

Table 6: Carrier-Grade NAT Translation Types

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

**Related Documentation** • [Junos OS Carrier-Grade NAT Implementation Overview on page 40](#)

## Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards

Network Address Translation (NAT) is a mechanism for translating IP addresses. NAT provides the technology used to support a wide range of networking goals, including:

- Concealing a set of host addresses on a private network behind a pool of public addresses.
- Providing a security measure to protect the host addresses from direct targeting in network attacks.
- Providing a tool set for coping with IPv4 address depletion and IPV6 transition issues.

The types of NAT supported by Junos OS are described in the following sections:

- [NAT Concept and Facilities Overview on page 44](#)
- [IPv4-to-IPv4 Basic NAT on page 45](#)
- [Static Destination NAT on page 45](#)
- [Twice NAT on page 46](#)
- [IPv6 NAT on page 48](#)
- [Application-Level Gateway \(ALG\) Support on page 48](#)
- [NAT-PT with DNS ALG on page 48](#)
- [Dynamic NAT on page 48](#)
- [Stateful NAT64 on page 49](#)
- [Dual-Stack Lite on page 49](#)

## NAT Concept and Facilities Overview

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

Junos OS 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 37](#).
- **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 39](#).
  - **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 37](#).
- **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 37](#).
- **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 129](#), [“NAT-PT with DNS ALG” on page 38](#), and [“Stateful NAT64” on page 39](#).
- **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 223](#).

Junos OS 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 41, 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 OS. 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*.



**NOTE:** With static NAT configured as basic NAT44 or destination NAT44 on MX Series routers with MS-MICs and MS-MPCs, the bits per second (bps) on the ingress side of the service interface might show a higher throughput than the egress side of the service interface under the Input bytes and Output bytes fields respectively in the output of the `show interface statistics` command. However, the packets per second (pps) values that are shown for the Input packets and Output packets fields in the output of the `show` command match correctly for the ingress and egress interfaces. The increased value for bps on the ingress side than the egress side is caused by the accounting of 16 bytes of the Juniper Forwarding Module (JFM) cookie. The tracing of packets indicate that the extra 16 bytes for packets arriving at the service PIC belong to the JFM cookie. Such packets are received from the ingress Packet Forwarding Engine (the Packet Forwarding Engine where packets enter the router) to the Packet Forwarding Engine or the lookup chip (LU) of the service PIC and these packets are sent out of the ms- logical interface to the PIC.

On MX Series routers with MS-MICs and MS-MPCs, it is noticed that the packet is looped if the source IP address of the packet is within the range of IP addresses configured in the NAT pool. The looping packet does not enable the timeout of the session to occur. With a NAT rule configured with NAPT44 and a source pool IP address specified, and with the NAT rule term match condition configured such that rule lookup fails, a spoofed packet with the source address that is the same as the NAT pool IP address, a session is created and its reverse flow destination IP address is the NAT pool IP address. This packet reaches the destination, which can reply a message. This response matches the reverse flow of the session and the packet is sent out of the PIC without translation. Because the destination address of packet is NAT pool IP address, the Packet Forwarding Engine sends the packet to the PIC again. It again matches the reverse flow and transmits the packet out of the PIC, resulting in a packet loop. We recommend that you manually clear the session and create a filter to block NAT pool IP spoofing in such a condition.

## 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 OS.



**NOTE:** Starting with Junos OS Release 15.1, the twice NAT functionality (**twice-basic-nat-44**, **twice-dynamic-nat-44**, and **twice-dynamic-napt-44** options) is supported on MX Series routers with MS-MPCs and MS-MICs.

In a network address translation (NAT) scenario with translation-type configured as twice-basic-nat-44, if the source-prefix and destination-prefix are configured, the destination-address is translated to the address that is configured as the source-prefix. Consider the following configuration with source prefix and destination prefix configured as the actions to be performed for NAT rule conditions:

```

nat {
  rule snat {
    match-direction output;
    term 1 {
      from {
        source-address {
          100.100.100.2/32;
        }
        destination-address {
          201.201.201.2/32;
        }
      }
      then {
        translated {
          source-prefix 101.101.101.2/32;
          destination-prefix 200.200.200.2/32;
          translation-type {
            twice-basic-nat-44;
          }
        }
      }
    }
  }
}

```

The following is a portion of the sample output of the **show services stateful-firewall conversations** command:

```

Number of initiators: 1, Number of responders: 1
Flow      100.100.100.2      -> 201.201.201.2      State   Dir      Frm count
ICMP      100.100.100.2      -> 201.201.201.2      Watch   0        9
  NAT source 100.100.100.2      -> 101.101.101.2
  NAT dest   201.201.201.2      -> 101.101.101.2

```

Here, the source address is translated as expected, whereas the destination address is translated to the same address as the source address. This problem occurs if the following conditions are satisfied:

- A NAT scenario with translation-type as twice-basic-nat-44

- Both the source-prefix and destination-prefix attributes are configured

## 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 OS.

## Application-Level Gateway (ALG) Support

Junos OS 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 55](#)

## 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 39](#).

**Figure 5: 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 39](#).

**Figure 6: 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 OS.

## Dual-Stack Lite

Dual-stack lite (DS-Lite) flow is shown in [Figure 4 on page 40](#).

**Figure 7: 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.



**NOTE:** In the output of the `show services sessions extensive` command, the Translation Type field displays the value as NAPT-44 for Endpoint Independent Filtering (EIF) flows. Also, the label, EIF, is displayed beside the translation type parameter to enable easy identification of EIF flows.

**Related  
Documentation**

- [Configuring Protocol Translation Between IPv6 and IPv4 Networks - NAT-PT on page 129](#)
- [Example: Configuring NAT-PT on page 136](#)
- [Configuring a DS-Lite Softwire Concentrator on page 237](#)

## CHAPTER 4

# NAT Configuration Overview

- [Network Address Translation Configuration Overview on page 51](#)
- [Configuring Source and Destination Addresses Network Address Translation Overview on page 52](#)
- [Configuring Pools of Addresses and Ports for Network Address Translation Overview on page 53](#)
- [Network Address Translation Rules Overview on page 55](#)
- [Configuring Service Sets for Network Address Translation on page 61](#)
- [Carrier-Grade NAT Implementation: Best Practices on page 62](#)

## 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 52](#).
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 53](#)
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 103](#).
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 55](#).
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 61](#).

### Related Documentation

- [Junos Address Aware Network Addressing Overview on page 35](#)
- [Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 41](#)

## 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**

- [Junos Address Aware Network Addressing Overview on page 35](#)

## Configuring Pools of Addresses and Ports for Network Address Translation Overview

- [Configuring NAT Pools on page 53](#)
- [Preserve Range and Preserve Parity on page 54](#)
- [Specifying Destination and Source Prefixes without Configuring a Pool on page 54](#)

### 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 55](#).

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 56](#)
- [Configuring Match Conditions in NAT Rules on page 56](#)
- [Configuring Actions in NAT Rules on page 57](#)
- [Configuring Translation Types on page 59](#)

## 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 9](#).
- 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 363.

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 53](#).
- 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 9.](#)

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 9](#)

---

## 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 63](#)
- [Do Not Use EIM with SIP on page 63](#)
- [Do Not Use EIM with HTTP, DNS, or When Not Needed on page 64](#)
- [Define PBA Blocks Based on User Profiles on page 65](#)



- [Do Not Change the PBA Configuration on Running Systems on page 65](#)
- [Do Not Allocate Excessively Large NAT Pools on page 66](#)
- [Configure the System Log for PBA Only When Needed on page 67](#)
- [Use Redundant Service PIC \(RSP\) Interfaces for Failover on page 69](#)
- [Contain the Effects of Missing IP Fragments on page 70](#)
- [Do Not Use Configurations Prone to Routing Loops on page 70](#)

## 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@host# 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@host# 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@host# 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@host# 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@host# 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@host# 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@host# 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 103](#)



## CHAPTER 5

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

- [Sample IPv6 Transition Scenarios on page 73](#)
- [Configuring Stateful NAT64 on page 75](#)

### 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 73](#)
- [Example 2: IPv4 Depletion with an IPv6 Access Network on page 74](#)
- [Example 3: IPv4 Depletion for Mobile Networks on page 75](#)

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

[Figure 8 on page 74](#) 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 8: 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 9 on page 74](#), the ISP network is IPv6-only.

Figure 9: 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 61](#).

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 6

# Hiding Private Networks Using Static Source NAT

- [Configuring Static Source Translation in IPv4 Networks on page 79](#)
- [Configuring Static Source Translation in IPv6 Networks on page 85](#)
- [Example: Configuring Basic NAT44 on page 89](#)
- [Example: Configuring NAT for Multicast Traffic on page 91](#)

## 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 79](#)
- [Configuring the Service Set for NAT on page 81](#)
- [Configuring Trace Options on page 83](#)
- [Sample Configuration - Static Source NAT Using a Static Pool With An Address Prefix And An Address Range on page 84](#)
- [Sample Configuration - Static Source Nat for One-To-One Mapping Between a Private Subnet and a Public Subnet on page 84](#)

## 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 355. 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
```

- 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
```

- 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
```

- 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;
  }
}
```

- 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
```

- 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 t1 {
        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 86](#)
- [Configuring the Service Set for NAT on page 87](#)
- [Configuring Trace Options on page 88](#)

## 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 89](#)
- [Overview on page 89](#)
- [Configuring Basic NAT44 on page 90](#)

### 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 10 on page 91 illustrates the network setup for the following configuration, which allows IP multicast traffic to be sent to the Multiservices PIC.

Figure 10: Configuring NAT for Multicast Traffic



- [Rendezvous Point Configuration on page 91](#)
- [Router 1 Configuration on page 94](#)

## 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 7

# Making Private Servers Available Using Static Destination NAT

- [Configuring Static Destination Address Translation in IPv4 Networks on page 97](#)

## 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;  
      }  
    }  
  }  
}
```

**Related  
Documentation**

- *Example: Configuring Static Destination Address Translation*





## CHAPTER 8

# 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 103](#)
- [Configuring Dynamic Source Address and Port Translation in IPv4 Networks on page 113](#)
- [Configuring Dynamic Source Address and Port Translation for IPv6 Networks on page 117](#)
- [Example: Configuring NAT with Port Translation on page 119](#)
- [Example: NAPT Configuration for the MS-MPC on page 120](#)
- [Example: Dynamic Source NAT as a Next-Hop Service on page 124](#)

### 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 104](#)
- [Sequential Allocation for NAPT on page 104](#)
- [Preserve Parity and Preserve Range for NAPT on page 105](#)
- [Address Pooling and Endpoint Independent Mapping for NAPT on page 105](#)
- [Port Block Allocation for NAPT on page 107](#)
- [Deterministic Port Block Allocation for NAPT on page 108](#)
- [Comparison of NAPT Implementation Methods on page 112](#)

## 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 105](#)
- [Endpoint Independent Mapping and Endpoint Independent Filtering on page 106](#)

---

### 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 107](#)
- [Interim Logging for Port Block Allocation on page 108](#)

---

### 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 NAPT

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 171](#).

- [Understanding Deterministic Port Block Allocation Algorithms on page 108](#)
- [Deterministic Port Block Allocation Algorithm Usage on page 109](#)
- [Deterministic Port Block Allocation Restrictions on page 111](#)

### 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— $(\text{Pu\_Offset} * \text{Port\_Range\_Per\_Pu\_IP}) + (\text{Pu\_Actual\_Port} - \text{Pu\_Port\_Start\_Port})$



**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.  $Pr\_Offset = 10.1.1.250 - 10.1.0.1 = 505$
2.  $Pu\_Port\_Offset = 505 * 249 = 125,745$
3.  $Rounded\_Port\_Range\_Per\_IP = \lceil (65,533/254) \rceil * 249 = 259 * 249 = 64,491$
4.  $Pu\_Prefix = 32.32.32.1 + \text{floor}(125,745 / 64,491) = 32.32.32.1 + 1 = 32.32.32.2$
5.  $Pu\_Start\_Port = 1,024 + (125,745 \% 64,491) = 62278$ 
  - 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.  $Pu\_Offset = Pu\_Prefix - Base\_Pu\_Prefix$
2.  $Pu\_Port\_Offset = (Pu\_Offset * Rounded\_Port\_Range\_Per\_IP) + (Pu\_Actual\_Port - Pu\_Port\_Range\_Start)$
3.  $Subscriber\_IP = Base\_Pr\_Prefix + \text{floor}(Pu\_Port\_Offset / Block\_Size)$

The reverse translation is determined as follows. Assume a flow returning to 32.32.32.2:62278.

1.  $Pu\_Offset = 32.32.32.2 - 32.32.32.1 = 1$
2.  $Pu\_Port\_Offset = (1 * 64,491) + (62,280 - 1024) = 125,747$
3.  $Subscriber\_IP = 10.1.0.1 + \text{floor}(125,747 / 249) = 10.1.0.1 + 505 = 10.1.1.250$



**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 7 on page 112](#).

Table 7: 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 <b>from</b> clause addresses should be same if the same deterministic NAT pool is used across multiple terms/rules. Only one <b>from</b> 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 <b>except</b> entries in the <b>from</b> 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 <b>address-allocation round-robin</b> 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 8 on page 112 provides a feature comparison of available NAPT implementation methods.

Table 8: 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	<b>active-block-timeout</b> feature	n/a

Table 8: Comparison of NAPT 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 113](#).

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 119](#)
- [Overview on page 119](#)
- [Configuring NAT with Port Translation on page 119](#)

### 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 120](#)
- [Overview on page 120](#)
- [Configuration on page 120](#)

### 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 NAPT<sup>44</sup> using the MS-MPC as a services interface card, perform these tasks:

- [Configuring Interfaces on page 121](#)
- [Configure an Application Set of Acceptable ALG traffic on page 122](#)
- [Configuring a Stateful Firewall Rule on page 122](#)

- [Configuring NAT Pool and Rule on page 123](#)
- [Configuring the Service Set on page 124](#)

**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**

- [Junos Address Aware Network Addressing Overview on page 35](#)

---

## 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 127](#)
- [Configuring Protocol Translation Between IPv6 and IPv4 Networks - NAT-PT on page 129](#)
- [Example: Configuring NAT-PT on page 136](#)

## ALGs Available by Default for Junos OS Address Aware NAT

The following Application Level Gateways (ALGs) listed in [Table 9 on page 127](#) 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 9: ALGs Available by Default**

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
Basic TCP ALG	yes	yes	<b>NOTE:</b> Specific Junos OS 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	<b>NOTE:</b> TCP tracker performs limited integrity and validation checks for UDP.

Table 9: 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"> <li>• junos-bootpc</li> <li>• junos-bootps</li> </ul>
DCE RPC Services	yes	yes	<ul style="list-style-type: none"> <li>• junos-dce-rpc-portmap</li> <li>• junos-dcerpc-endpoint-mapper-service</li> <li>• junos-dcerpc-msexchange-directory-nsp</li> <li>• junos-dcerpc-msexchange-directory-rfr</li> <li>• junos-dcerpc-msexchange-information-store</li> </ul>
DNS	yes	yes	<ul style="list-style-type: none"> <li>• junos-dns-tcp</li> <li>• junos-dns-udp</li> </ul>
FTP	yes	yes	<ul style="list-style-type: none"> <li>• junos-ftp</li> </ul>
H323	yes	no	<ul style="list-style-type: none"> <li>• junos-h323</li> </ul>
ICMP	yes	yes	<ul style="list-style-type: none"> <li>• junos-icmp-all</li> <li>• junos-icmp-ping</li> </ul>
IIOp	yes	no	<ul style="list-style-type: none"> <li>• junos-iiop-java</li> <li>• junos-iiop-orbix</li> </ul>
IP	yes	The TCP tracker, available by default on these platforms, performs limited integrity and validation checks.	<ul style="list-style-type: none"> <li>• junos-ip</li> </ul>
NETBIOS	yes	no	<ul style="list-style-type: none"> <li>• junos-netbios-datagram</li> <li>• junos-netbios-name-tcp</li> <li>• junos-netbios-name-udp</li> <li>• junos-netbios-session</li> </ul>
NETSHOW	yes	no	<ul style="list-style-type: none"> <li>• junos-netshow</li> </ul>
PPTP	yes	yes	<ul style="list-style-type: none"> <li>• junos-pptp</li> </ul>
REALAUDIO	yes	no	<ul style="list-style-type: none"> <li>• junos-realaudio</li> </ul>
Sun RPC and RPC Port Map Services	yes	yes	<ul style="list-style-type: none"> <li>• junos-rpc-portmap-tcp</li> <li>• junos-rpc-portmap-udp</li> </ul>
RTSP	yes	yes	<ul style="list-style-type: none"> <li>• junos-rtsp</li> </ul>

Table 9: 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"> <li>• <b>junos-sip</b></li> </ul> <p>The SIP <b>callid</b> is <i>not</i> translated in <b>register</b> messages.</p> <p><b>NOTE:</b> SIP sessions are limited to 12 hours (720 minutes) for NAT processing on the MS-MIC and MS-MPC interface cards. SIP sessions on the MS-DPC have no time limits.</p>
SNMP	yes	No	<ul style="list-style-type: none"> <li>• <b>junos-snmp-get</b></li> <li>• <b>junos-snmp-get-next</b></li> <li>• <b>junos-snmp-response junos-snmp-trap</b></li> </ul>
SQLNET	yes	yes	<ul style="list-style-type: none"> <li>• <b>junos-sqlnet</b></li> </ul>
TFTP	yes	yes	<ul style="list-style-type: none"> <li>• <b>junos-tftp</b></li> </ul>
Traceroute	yes	yes	<ul style="list-style-type: none"> <li>• <b>junos-traceroute</b></li> </ul>
Unix Remote Shell Service	yes	Yes	<ul style="list-style-type: none"> <li>• <b>junos-rsh</b></li> </ul>
WINFrame	yes	No	<ul style="list-style-type: none"> <li>• <b>junos-citrix-winframe</b></li> <li>• <b>junos-citrix-winframe-udp</b></li> </ul>
TALK-UDP	No	Yes	<ul style="list-style-type: none"> <li>• <b>junos-talk-udp</b></li> </ul>
MS RPC	No	Yes	<ul style="list-style-type: none"> <li>• <b>junos-rpc-portmap-tcp</b></li> <li>• <b>junos-rpc-portmap-udp</b></li> <li>• <b>junos-rpc-services-tcp</b></li> <li>• <b>junos-rpc-services-udp</b></li> </ul>

**Related Documentation** • [ALG Descriptions on page 299](#)

## 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 130](#)
- [Configuring the NAT Pool and NAT Rule on page 130](#)
- [Configuring the Service Set for NAT on page 133](#)
- [Configuring Trace Options on page 134](#)

## 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 136](#)
- [Overview and Topology on page 136](#)
- [Configuration of NAT-PT with DNS ALGs on page 138](#)

## 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 11: 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 138](#)
- [Configuring the NAT Pools on page 139](#)
- [Configuring the DNS Server Session: First NAT Rule on page 140](#)
- [Configuring the HTTP Session: Second NAT Rule on page 143](#)
- [Configuring the Service Set on page 145](#)
- [Configuring the Stateful Firewall Rule on page 147](#)
- [Configuring Interfaces on page 148](#)

### 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 130](#). 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 139 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](http://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](http://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 147](#).

```
[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 148](#).

**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 148](#).  

```
[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 on page 43](#)
  - [Configuring Protocol Translation Between IPv6 and IPv4 Networks - NAT-PT on page 129](#)
  - [Configuring Service Sets to be Applied to Services Interfaces on page 9](#)
  - [Example: Configuring Layer 3 Services and the Services SDK on Two PICs on page 367](#)
  - [dns-alg-prefix on page 759](#)
  - [dns-alg-pool on page 758](#)





## CHAPTER 10

# Reducing Traffic and Bandwidth Requirements Using Port Control Protocol

- [Port Control Protocol Overview on page 151](#)
- [Configuring Port Control Protocol on page 153](#)
- [Example: Configuring Port Control Protocol with NAPT44 on page 156](#)

### Port Control Protocol Overview

---

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 keepalive traffic. PCP is designed to be implemented in the context of both Carrier-Grade NATs (CGNs) and small NATs (for example, residential NATs). PCP enables 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 ISP. PCP enables 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 a mapping for incoming connections is created, remote computers must be informed about the IP address and port for the incoming connection. This is usually done in an application-specific manner.

Junos OS supports PCP version 2 and version 1.

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 sessions are not being timed out by a NAT. These applications can reduce the frequency of such NAT keepalive messages by using PCP to learn and influence the NAT mapping lifetime. Using PCP helps reduce bandwidth on the subscriber's access network, traffic to the server, and battery consumption on mobile devices.

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 user equipment, as shown in [Figure 12 on page 152](#).

Figure 12: 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 13 on page 152](#)

Figure 13: PCP with DS-Lite Plain Mode



**NOTE:** Junos OS does not support deterministic port block allocation for PCP-originated traffic.

## Port Control Protocol Version 2

Starting with Junos OS Release 15.1, Port Control Protocol (PCP) version 2 is supported, which is in compliance with RFC 6887. 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 version 2 supports nonce authentication. PCP allows applications to create mappings from an external IP address and port to an internal IP address and port. A nonce payload prevents a replay attack. A nonce payload is sent by default unless it is explicitly disabled. With nonce payload, the

device expects Online Certificate Status Protocol (OCSP) responses to contain a nonce payload, otherwise the revocation check will fail. If OCSP responders are not capable of responding with a nonce payload, disable this option.

Client nonce verification for version 2 map requests (for refresh or delete) requires that the nonce received in the original map request that causes the PCP mapping to be created is preserved. The version of the initial request that enables the mapping to be created is also preserved. This behavior of saving the nonce and version parameters denotes that 13 bytes per PCP mapping are used. This slight increase in storage space is not significant when matched with the current memory usage of a system for a single requested mapping (taking into account the endpoint-independent mapping (EIM) and endpoint-independent filtering (EIF) that are created along with it). In a customer deployment, PCP causes EIM and EIF mappings to represent a fraction of all such mappings.

Until Junos Release 15.1, services PICs support PCP servers on Juniper Networks routers in accordance with PCP draft version 22 with version 1 message encoding. With PCP being refined from the draft version as defined in *Port Control Protocol (PCP) draft-ietf-pcp-base-22 (July 2012 expiration)* to a finalized, standard version as defined in RFC 6887 -- Port Control Protocol (PCP), the message encoding changed to version 2 with the addition of a random nonce payload to authenticate peer and map requests as necessary. Version 1 does not decode messages compliant with version 2 format and nonce authentication is not supported. In a real-world network environment, with customer premises equipment (CPE) devices increasingly supporting version 2 only, it is required to parse and send version 2 messages. Backward compatibility with version 1-supporting CPE devices is maintained (version negotiation is part of the standard) and authenticates request nonce payload packets when v2 messages are in use.

The output of the **show services pcp statistics** command contains the PCP unsupported version field, which is incremented to indicate whenever the version is not 1 or 2. A new field, PCP request nonce does not match existing mapping, is introduced to indicate the number of PCP version 2 requests that were ignored because the nonce payload did not match the one recorded in the mapping (authentication failed). If version 2 is in use, the client nonce is used for authentication.

**Related Documentation**

- [Configuring Port Control Protocol on page 153](#)

---

## Configuring Port Control Protocol

---

This topic describes the following configuration tasks:

- [Configuring PCP Server Options on page 153](#)
- [Configuring a PCP Rule on page 155](#)
- [Configuring a Service Set to Apply PCP on page 155](#)
- [SYSLOG Message Configuration on page 156](#)

### 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
```

4. 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
```

5. 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
```

6. (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
```

7. (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.

---

8. (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 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

1. 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 156](#)
- [Overview on page 156](#)
- [PCP Configuration on page 157](#)

### 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 14 on page 157](#) shows the basic topology for this example.

Figure 14: PCP with NAT44



## 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-minimum 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** user@host# 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@host# 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@host# set nat-rules pcp-rule
```

3. Identify the PCP rule associated with the service set.

```
[edit services service-set sset_0]
user@host# set pcp-rules r1
```

4. Identify the service interface associated with the service set.

```
[edit services service-set sset_0]
user@host# 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;
}
```



## CHAPTER 11

# Automatically Assigning Ports Using Port Block Allocation

- [Secured Port Block Allocation for NATP on page 163](#)
- [Interim Logging for Port Block Allocation on page 163](#)
- [Guidelines for Configuring Interim Logging for Secured Port Block Allocation on page 164](#)
- [Secured Port Block Allocation for NAT44 and NAT64 Events on MS-MPCs and MS-MICs Overview on page 167](#)
- [Guidelines for Configuring Secured Port Block Allocation for MS-MPCs and MS-MICs on page 167](#)
- [Configuring Secured Port Block Allocation on page 169](#)
- [Configuring Deterministic Port Block Allocation on page 171](#)

## Secured Port Block Allocation for NATP

---

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 169](#)

## 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 traffic flows that last for a long duration, the ALLOC and RELEASE syslogs can be archived in different records in the customer deployment. Interim logs enable you to identify the currently used port blocks, thereby eliminating the need to search and analyze archived logs to identify the internal host that is using the external IP address and port (from the port block recorded in the log message). Depending on your network topology, you can set the interval for the port block allocation logs based on the period of the archive so that at least one log per port block (for an active flow) in each archive is present. To configure the interim logging interval at the MS-PIC level, which applies to all the NAT pools on that ms- interface, include the **pba-interim-logging-interval seconds** statement at the **[edit interfaces ms-fpc/pic/port services-options]** hierarchy level. To configure the interim logging interval at a NAT pool level, include the **interim-logging-interval seconds** statement at the **[edit services nat pool pool-name port secured-port-block-allocation]** hierarchy level. You can specify a value from 0 through 86400 seconds for the interim logging frequency. This capability is supported on MX Series routers with MS-MPCs and MS-MICs.

**Related  
Documentation**

- [Configuring NAT Session Logs on page 219](#)
- [Secured Port Block Allocation for NAPT on page 163](#)

---

## Guidelines for Configuring Interim Logging for Secured Port Block Allocation

---

Observe the following guidelines when you configure the interim logging interval for secured port block allocation:

- Interim logging is enabled only when the interim logging functionality is configured. The **pba-interim-logging-interval** statement that you can configure at the **[edit interfaces ms-fpc/pic/port services-options]** hierarchy level of an ms-interface is provided for backward compatibility (on the uKernel, this feature is configured under service-options of the services PIC (sp-) interface). The **interim-logging-interval** statement that is available for configuration for each NAT pool extends the set of parameters that you can already configure for PBA.
- If you configure the interim logging capability to be applicable to all PBA pools residing on that particular MS-PIC and the interim logging capability for a specific PBA pool, the NAT pool-specific interval takes precedence over the MS-PIC-specific interval. For port blocks allocated from other PBA pools for which interim logging interval at the NAT pool-level is not configured, the logging interval value as configured at the ms-interface-level applies.
- The default value is zero, which denotes no interim logging message is generated.
- Interim logs are sent any time after the configured period of time in seconds. The time-difference is not fixed between the logging intervals of two logs.

- Interim logs are generated for port blocks (both active and inactive) that contain at least one port in use by a flow which has traffic. No timer controls run on the port blocks to generate the logs. When a packet is received on a flow, the validation is performed to generate an interim log. If the conditions are satisfied, an interim log is generated for that port block. Interim logs are not generated for deleted port blocks.
- The interim log contains the timestamp of the port block creation in hexadecimal format (when local time is set, the hexadecimal value provides the time in UTC format).
- The conversion of the timestamp to UTC format can be performed in the external syslog server as necessary.
- In certain scenarios, it is possible that the timestamp in hexadecimal value and the actual timestamp in ALLOC messages differ by a couple of seconds. This behavior occurs because the syslog mechanism contains a slight difference when it reads the time (as seen in PORT\_BLOCK\_ALLOC syslog) and the time at which NAT application reads the time (to update the ALLOC time in the subscriber context). The interim system log displays the ALLOC time retrieved from the subscriber context.
- Because these logs are generated on CPU computation and in the fast path, a slight impact might be observed with fast path performance only when a generation of the log occurs.
- Port block creation timestamp in hexadecimal is saved in the JSERVICES\_NAT\_PORT\_BLOCK\_RELEASE message, even if interim logging is not present.
- If you define the logging interval when traffic flow is in progress, this functionality takes effect on existing and new flows. You need not reboot the MIC or activate and deactivate the service set.
- If the flows or subscribers are timing out, it denotes that no new packets or traffic flows are seen for this 5-tuple data or for that particular subscriber. In such a case, interim logs are not generated.
- If the interim-logging interval is lower than the inactivity-timeout of the flow, interim logs are not observed when the flow is timing out and the interim-logging interval has elapsed. If the interim-logging interval is lower than the subscriber-timeout value, interim logs are not observed when the subscriber is timing out and the interim-logging interval has elapsed. For example, if the inactivity-timeout is configured to 2500 seconds and the interim-logging is configured as 1800 seconds, when the flow is timing out, there is a point in time when 1800 seconds has elapsed since the last packet was seen on this flow and no interim log is generated in this case.
- The interim logs are recorded for those pools that have PBA configured. If pools exist without the PBA configuration present on the service network processing unit (NPU), interim logs are not saved even if you enable the interim logging functionality.
- You can configure only a range of values for the interval at which the logs need to be generated, such as 0, [1800, 86400].
- You can enable the generation of syslogs by using the *syslog* statement at the **[edit system]** and **[edit services service-sets service-set name nat rule rule-name term term-name then]** hierarchy levels that contain the NAT rules with PBA pools. Interim logs are not triggered if the recording of syslogs are not enabled on the system.

- We recommend that you configure the interim-logging interval to be higher than the inactivity timeout period for established flows. Also, we recommend that you configure the interim-logging interval to be higher than the subscriber-timeout value. When endpoint-independent mapping (EIM) is configured, the interim-logging interval must be higher than the sum of the address pooling paired (APP) timeout and EIM timeout values.
- Transmission of logs occurs in clear-text format similar to other log messages that the services PICs do not encrypt. It is assumed that the transport of logs and the positioning of the log collector are within a secured realm. Because the messages do not contain sensitive details such as username or passwords, the messages do not cause any security or reliability risks. Increased generation of log messages does not cause a possibility of a flood of logs because the frequency of logging can be configured, depending on the network topology, traffic levels, and your monitoring needs.
- The logs for PBA in the microkernel start with the prefix of ASP\_\*. These logs have been modified to start with the prefix of JSERVICES\_\*. The following are examples of system logs for PBA in the microkernel and with the Junos OS Extension-Provider packages installed and configured on the device.

**Microkernel:** 1970-01-01 00:32:36 {nat64}[FWNAT]:ASP\_NAT\_PORT\_BLOCK\_ACTIVE: 2001:0:0:0:0:0:2 -> 1.1.1.1:1050-1091 0x6f

**Junos OS Extension-Provider (eJunos):** 1970-01-01 00:32:36 {nat64}[FWNAT]:JSERVICES\_NAT\_PORT\_BLOCK\_ACTIVE: 2001:0:0:0:0:0:2 -> 1.1.1.1:1050-1091 0x6f

- Also, you can specify the interim logging interval per NAT pool instead of a global configuration per MS-PIC, based on whether you want the syslog settings to apply to all the NAT pools on a device or for a particular NAT pool. For NAT, the member interfaces must have the jservices-nat package configured. The JSERVICES\_NAT\_PORT\_BLOCK\_ACTIVE system logging message is generated when you configure interim logging for PBA. The following sample logs denote the log messages generated with the interim interval set as 1800 seconds. You can notice that the timestamp between consecutive interim logs is more than 1800 seconds.

```
1970-01-01 00:01:51 [FWNAT]:JSERVICES_NAT_PORT_BLOCK_ALLOC: 2001:0:0:0:0:0:2
-> 1.1.1.1:1050-1091
1970-01-01 00:32:36 {nat64}[FWNAT]:JSERVICES_NAT_PORT_BLOCK_ACTIVE:
2001:0:0:0:0:0:2 -> 1.1.1.1:1050-1091 0x6f
1970-01-01 01:03:20 {nat64}[FWNAT]:JSERVICES_NAT_PORT_BLOCK_ACTIVE:
2001:0:0:0:0:0:2 -> 1.1.1.1:1050-1091 0x6f
1970-01-01 01:34:04 {nat64}[FWNAT]:JSERVICES_NAT_PORT_BLOCK_ACTIVE:
2001:0:0:0:0:0:2 -> 1.1.1.1:1050-1091 0x6f
1970-01-01 02:04:48 {nat64}[FWNAT]:JSERVICES_NAT_PORT_BLOCK_ACTIVE:
2001:0:0:0:0:0:2 -> 1.1.1.1:1050-1091 0x6f
```

#### Related Documentation

- [Configuring NAT Session Logs on page 219](#)
- [Secured Port Block Allocation for NAPT on page 163](#)



## Secured Port Block Allocation for NAT44 and NAT64 Events on MS-MPCs and MS-MICs Overview

Starting with Junos OS Release 15.1, in an environment in which Junos Address Aware (carrier-grade NAT) is employed, service providers or carrier operators can monitor and track the consumption of resources and types of services being utilized by subscribers or users in an easier and effective manner by using system logging messages recorded for the allocation of ports to clients. By using IP addresses in RADIUS or DHCP logs, evaluation of the logs is performed to analyze and determine the services usage and bandwidth consumption by subscribers. With carrier-grade NAT, because IP addresses are shared by multiple subscribers, examining logs to track the IP addresses and ports that are part of the system logs might be time-consuming and difficult. Also, because ports are allocated and released at frequent intervals depending on the logging-in and closure of subscriber sessions, a large number of logs are triggered for every port allocation and deallocation. As a result, excessive syslogs render it cumbersome to archive and correlate the logs to identify a subscriber. You can now allocate ports in blocks, which reduces the amount of syslogs considerably.

You can use the secured port block allocation (PBA) mechanism to allocate ports in blocks for NAT44 (translation of an IPv4 address to an IPv4 address) and NAT64 (translation of an IPv6 address to an IPv4 address) types. By using secured PBA, the port usage might be a little inefficient, depending on traffic patterns. Allocation of port blocks for NAT44 and NAT64 types is supported on MX Series routers with MS-MPCs and MS-MICs.

Port Block Allocation (PBA) ensures that when a subscriber requires a port to be assigned for the first time, a block of ports are allocated to the particular user. Here, a subscriber is defined uniquely as a private IP address and service set ID. Because the subscriber has a block of ports assigned to it, all subsequent requests from this subscriber use ports from the assigned block. A new port block is allocated when the current active block is exhausted, or after the active port block timeout interval has expired. The maximum number of blocks allocated to a user is defined by the **max-blocks-per-address** option that you can specify at the **[edit services nat pool pool-name port secured-port-block-allocation]** hierarchy level. This behavior of allocation of NAT ports in blocks is different from the traditional NAT utility where the request for a port allocates a single port and not a group of ports in a block. Deterministic PBA is not supported for MS-MPCs and MS-MICs.

## Guidelines for Configuring Secured Port Block Allocation for MS-MPCs and MS-MICs

Keep the following points in mind when you configure secured PBA for MS-MPCs and MS-MICs:

- Block size is not configurable at the NAT rule level.
- Increase in setup rate of sessions is not impacted when you configure secured PBA.
- If a block of a particular size is not available, an out-of-ports message is displayed and smaller-sized blocks are not allocated alternatively in such a scenario.

- Addresses in the pool using port-block-allocation method cannot be used in any other pool.
- Port range in the NAT pool must be contiguous.
- Preserve parity (Allocate ports with same parity as the original port) is not supported with block-allocation of ports.
- The limitation on the number of open sessions when the specified threshold is reached (for intrusion detection services) and the maximum number of blocks that can be allocated to a user address that is configured for secured PBA are independent functionalities.
- The functionality to preserve privileged port range after translation is not supported. The blocks are assigned from unprivileged port range (1024-65535). For ports in privileged range, port block allocation method is not applicable.
- Port usage efficiency is lower when port-block allocation is enabled. PBA does not use ports from 0-1023 of a NAT IP address.
- If you configure the automatic port assignment method, which enables sequential assignment of ports, the port range from 1024 through 65535 is available for allocation to users.
- Port blocks can start at any start port that you can configure.
- The number of ports used is dependent on the block size and the rest of the ports are not be used.
- An overloaded pool, which indicates an address pool that can be used if the source pool becomes exhausted, is not supported with secured PBA.
- NAT IP addresses of PBA pool must not overlap with any other pool. Although a validation is not performed to identify whether any overlapping pools exist, you must ensure that the addresses of a pool that is used for PBA are not used in other pools. This condition is because some of the users require the overload pool to use the same IP addresses as that of NAT IP addresses, but a different port range of PBA pool to support the address pooling paired (APP) functionality.
- The block-size is fixed per NAT pool and is configurable at the NAT pool level. Multiple port blocks can be allocated to a private IP address.
- You can configure the maximum number of blocks per pool per subscriber by including the **max-blocks-per-user max-blocks** statement at the **[edit services nat pool pool-name port secured-port-block-allocation]** hierarchy level. If a subscriber matches two pools, that particular user can be allocated a maximum of port blocks that equals the sum of the maximum number of port blocks for each pool for that subscriber. New requests for NAT ports arrive from the current active block only.
- Ports can be allocated randomly from the current active block, which specifies whether ports should be allocated sequentially or randomly within the port block.
- A block is active for a timeout interval that you can define by including the **active-block-timeout timeout-seconds** at the **[edit services nat pool pool-name port secured-port-block-allocation]** hierarchy level. After the timeout period, a new block is allocated even if ports are available in the active block. The default timeout of an

active block is 120 seconds. When you configure it as 0 (infinite), the active block transitions to inactive only when it runs out of ports and a new block is allocated.

- If the maximum number of blocks of blocks is exceeded, and a new request is received, the active block is moved to a block that contains available ports. Any non-active block without any ports in use is freed to NAT pool.
- In addition to tracking port blocks assigned to each private IP address, actual ports in use are also computed and maintained. This metric is used to calculate port usage efficiency.
- A syslog message is generated for each block allocation and release. The format of the message is similar to the messages recorded for individual port allocation and release.
- Session setup rate is the same or slightly improved than the existing non-block allocation setup rate. NAT pool using block-port allocation method can have partial port ranges. If the address is used for port forwarding, those ports can be removed from the pool port range. You can configure partial port ranges by using the **port range low minimum-value high maximum-value random-allocation** statement at the **[edit services nat pool nat-pool-name]** hierarchy level. Port block allocation works in the same manner as NAT44 for TCP, UDP, and ICMP traffic.
- Randomness can be achieved by allocating ports randomly within the block and changing active block periodically. The block of ports do not contain random ports (ports within the block are sequential). This capability is supported with aggregated multiservices (ams) interfaces.
- The starting port number is calculated differently in the microkernel and in Junos OS Extension-Provider packages. In the microkernel, the starting or first port is the nearest multiple of the block size after 1023. In that implementation, more ports are wasted because ports are wasted at the beginning and the end of the port range depending on the block size. In Junos OS Extension-Provider packages, the start port of a block is not restricted to a multiple of the block size. The start port can start at the lower boundary of the range of the port configured.

#### Related Documentation

- [Configuring NAT Session Logs on page 219](#)
- [Secured Port Block Allocation for NAT on page 163](#)

## 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 51](#)

## 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 51](#)

## CHAPTER 12

# Connecting Specific Ports and Addresses Using Port Forwarding

- [Configuring Port Forwarding for Static Destination Address Translation on page 173](#)
- [Configuring Port Forwarding Without Destination Address Translation on page 176](#)
- [Example: Configuring Port Forwarding with Twice NAT on page 177](#)

### 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-name 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 177.
- 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 97](#)

## 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 173](#)



# Allocating a Few Public Addresses to Many Private Hosts Using Dynamic NAT

- [Configuring Dynamic Address-Only Source Translation in IPv4 Networks on page 181](#)
- [Example: Dynamic Source NAT as a Next-Hop Service on page 185](#)
- [Example: Assigning Addresses from a Dynamic Pool for Static Use on page 187](#)

## 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 189](#)
- [Example: Configuring Inline Network Address Translation - Interface-Service Service Set on page 191](#)

## 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:

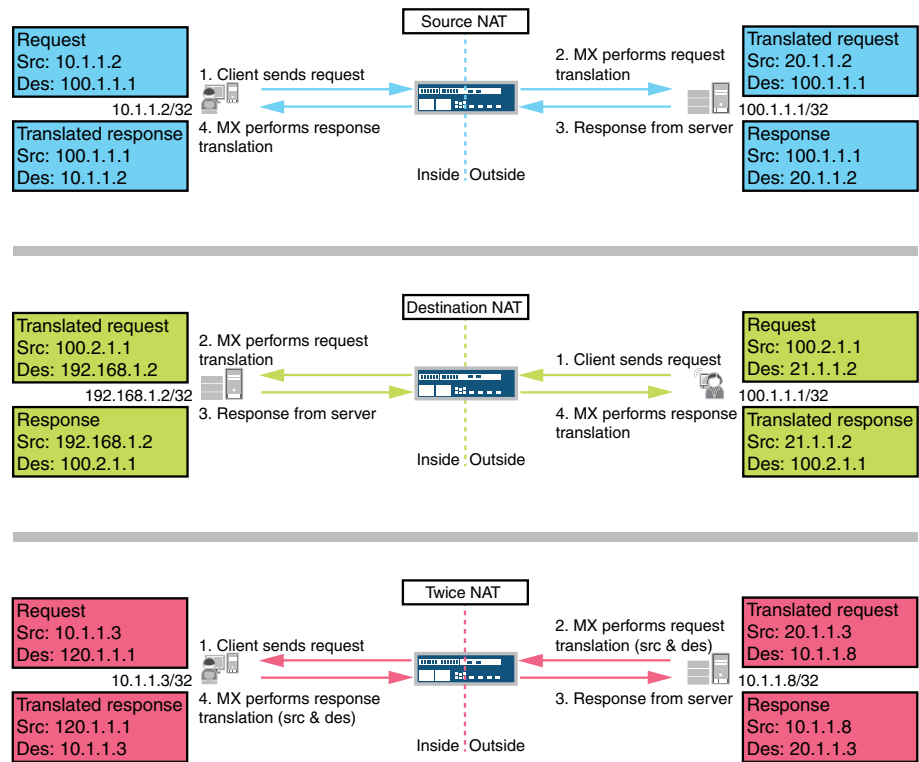
- 1:1 static address mapping
- Bidirectional mapping - source NAT for outbound traffic and destination NAT for inbound traffic
- No limit on number of flows
- Support for Source, destination, and twice NAT, as shown in [Figure 15 on page 190](#)



**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 15: 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 51](#)
- [Example: Configuring Inline Network Address Translation - Interface-Service Service Set on page 191](#)
- [Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 41](#)



## Example: Configuring Inline Network Address Translation - Interface-Service Service Set

---

- [Requirements on page 191](#)
- [Overview on page 191](#)
- [Configuration on page 193](#)

### 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 16 on page 192](#)

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 16 on page 192

**Figure 16: 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 193](#)
- [Configuring Bandwidth for the Service Inline \(si-\) Interface on page 195](#)
- [Configuring NAT Pool and Rule on page 196](#)
- [Configuring the Service Set on page 197](#)

### 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 192.1.68.1/24
set services nat pool p2 address 192.1.68.2/24
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 pool p1 address 192.168.1.0/24  
user@host# set pool p2 address 192.168.2.0/24
```
3. Configure NAT rule for source pool p1.  

```
[edit services nat]  
user@host# set rule r1 match-direction input  
user@host# set rule r1 term t1 from source-address 10.1.1.0/24  
user@host# set rule r1 term t1 then translated source-pool p1 translation-type  
basic-nat44
```
4. Configure NAT rule for source pool p2.  

```
[edit services nat]  
user@host# set rule r2 match-direction input  
user@host# set rule r2 term t1 from source-address 10.1.3.0/24  
user@host# set rule r2 term t1 then translated source-pool p2 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 189](#)



## CHAPTER 15

# Removing Address Dependency Using Network Prefix Translation for IPv6 Traffic

- [Attaining Address Independence by Eliminating the Need for Advertising Internal Prefixes Using NPTv6 on page 199](#)
- [Guidelines for Configuring Stateless Source Network Prefix Translation on page 201](#)
- [Interoperation of Functionalities with Network Prefix Translation for IPv6 on page 202](#)
- [Working of NPTv6 with Interface-Style and Next Hop-Style Service Sets on page 204](#)
- [Example: Achieving Address Independence By Configuring Stateless Network Prefix Translation in IPv6 Networks by Using Interface-Style Service Sets on page 205](#)
- [Example: Achieving Address Independence By Configuring Stateless Network Prefix Translation in IPv6 Networks by Using Next-Hop -Style Service Sets on page 211](#)

## Attaining Address Independence by Eliminating the Need for Advertising Internal Prefixes Using NPTv6

---

Starting with Junos OS Release 15.1, you can configure stateless translation of source address prefixes in IPv6 networks (IPv6 to IPv6). This capability is supported on MX Series routers with MPCs where inline NAT is supported. To configure stateless network prefix translation for IPv6 packets (NPTv6), include the **translation-type nptv6** statement at the **[edit services nat rule *rule-name* term *term-name* then translated]** hierarchy level. The NPTv6 translator translates the source address prefix in such a way that the transport layer checksum of the packet does not need to be recomputed. NPTv6 defines a stateless method of IPv6 network prefix translation between internal and external networks. NPTv6 does not maintain the state for each node or each flow in the translator. You can use the **show services nat mappings nptv6 (internal | external)** command to view the NAT mappings for NPTv6 for internal and external addresses respectively. You can also use the **show services inline nat statistics** and **show services inline nat pool** commands to display information about inline NAT with NPTv6 configured.

Network prefix translation for IPv6 (NPTv6) defines a stateless way of IPv6 address prefix translation between internal and external networks. NPTv6 does not maintain the state for each node or each flow in the translator. Maintenance of mapping state is not required for the address mapping of inbound or outbound packets. A stateless, transport-agnostic IPv6-to-IPv6 NPTv6 function offers the advantage of address-independence associated with IPv4-to-IPv4 NAT (NAPT44) and provides a 1:1

relationship between addresses in the *inside* and *outside* prefixes, thereby preserving end-to-end reachability at the network layer.

IPv6 address independence contains the following characteristics:

- In edge networks, it is not necessary to renumber the IPv6 addresses used inside the local network for interfaces, access lists, and system logging messages in the following scenarios:
  - If the global prefixes assigned for use by the edge network are changed
  - If the IPv6 addresses are used inside the edge network or within other upstream networks (such as multihomed devices) when a site adds, drops, or changes upstream networks

It is not necessary for an organization to reconfigure the upstream network to route its internal IPv6 prefixes or for it to advertise prefixes derived from other upstream networks into it.

- In upstream networks, IPv6 addresses used by the edge network always contain a provider-allocated prefix, thereby removing the requirement and the processing for ingress filtering and the advertisement of customer-specific prefixes.

NPTv6 is designed to provide address independence to the edge networks to achieve internal address stability, regardless of its upstream service provider networks. However, using provider-independent addresses without translation might be very expensive because the routing table enumerates the edge networks, instead of enumerating the transit domain that provides the service to the edge networks. This phenomenon can cause a massive and unmanageable route table. NPTv6 is a mechanism that effectively and cohesively provides address independence without advertising an internal network prefix to external networks. In contrast, the main objective of network address port translation (NAPT) for IPv4 (NAPT44) is to solve IPv4 address depletion, although it brings the same benefit of address independence. NAPT for IPv6, specifically NAPT66, is already supported in microkernel. However, similar to NAPT44, NAPT66 requires flow-state information to be preserved. NPTv6 provides a simple and streamlined technique to avoid as many of the limitations associated with NAPT66 as possible. It is defined to include a two-way, checksum-neutral, and an algorithmic translation function.

NPTv6 does not maintain state information for a node, flow, or a connection in the translator. Internal to external and external to internal packets are translated algorithmically using information present in the IPv6 header. As a result of its stateless nature, a reset or brief outage of an NPTv6 translator does not disrupt connections that traverse the translation function, and if multiple NPTv6 translators exist between the same two networks, the load can shift or be dynamically shared among them. Also, unlike NAPT44, because the mapping can be done in either direction, the translator does not interfere with the inbound connection establishment. Instead, a firewall can be used in conjunction with an NPTv6 translator. This behavior offers the network administrator more flexibility to specify security policy than that can be achieved with a traditional NAT.

Another advantage of NPTv6 is checksum-neutral translations. The translator does not need to rewrite the transport header for updating the checksum and does not perform port mapping. As a result, to deploy new transport layer protocols, you do not need to

modify the translator. Because the transport layer is not modified, the algorithm does not interfere with encryption of the IP payload. Although NPTv6 compares favorably to NAT44 or NAT66 in several ways, it does not eliminate all of the architectural problems. Because NPTv6 modifies the IP headers of packets, it is not compatible with security mechanisms such as the IPsec authentication header. The use of separate internal and external prefixes creates complexity for Domain Name System (DNS) deployment. Also, those applications that require application layer gateways (ALGs) to work correctly through NAT44 or NAT66 devices might require similar ALGs to work through NPTv6 translators. Because NPTv6 does not maintain connection states, the failure of the translator does not impact the non-transmit power control (TPC) traffic through the server. TCP connections can be interrupted because of the change in the source IP address of a connection. Connections might be timed out and then reestablished in this case.

NPTv6 is available as an MS-DPC implementation and uses inline NAT. Inline NAT uses the capabilities of the Modular Port Concentrator (MPC) line card, eliminating the need for a MultiServices DPC (MS-DPC) for NAT. 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 service sets and next-hop service sets used for NAT. The **si-** interface serves as a *virtual service PIC*.

**Related** •  
**Documentation**

## Guidelines for Configuring Stateless Source Network Prefix Translation

Keep the following points in mind when you configure stateless translation of source IPv6 prefixes:

This topic contains the following sections that describe the working behavior of different functionalities with stateless source IPv6 prefix translation and the various system conditions:

- Graceful Routing Engine switchover (GRES) support is the same as for NAT44.
- Unified ISSU and nonstop software upgrade (NSSU) are not supported.
- NPTv6 deployment enables direct inbound connections to internal nodes from external networks. This mechanism causes slight vulnerability because it opens the internal nodes to attacks from outside. The stateless translation of NPTv6 makes it difficult to trace external connection requests, based on connection states. This behavior enables NAT44 networks to be well-protected against external attacks. The best option to secure an NPTv6 translator is to add a firewall above the NPTv6 translator.
- A 6rd software concentrator interoperates with NPTv6. All other mechanisms that do not require the application layer gateway (ALG) to change the source IP address in the payload are supported. TCP, UDP, ICMP, SSH, and Telnet are supported by the NPTv6 translator. FTP and Session Initiation Protocol (SIP) that require the ALG to change the source IP address in the payload are not supported.
- The NPTv6 pools are allocated in the external data memory. The pool data structure consists of the address prefix, prefix length, and the checksum. The size of each record

is of 192 bits. For every pool, a denat pool is allocated automatically. The size of the denat pool is 192 bits. There is a total allocation of 8000 64-bit entries for NAT-processed and untranslated NPTv6 pools. This allocation comes from the 64,000 entries allocated for the inline services (JNH\_APP\_INLINE\_SVCS).

- Chaining of inline services for interoperation of 6rd with NPTv6 is not supported.
- You need to configure a source pool and specify the **from** (source) address, while configuring NPTv6.
- The external and internal prefix lengths must be greater than or equal to /16 subnet mask and less than or equal to /112 subnet mask.
- Two different internal prefixes cannot be translated to the same external prefix.
- NPTv6 cannot be applied to IPSec and Internet Key Exchange (IKE) packets. The NPTv6 translator is bypassed in this case.
- Because the translation is of one IPv6 address prefix, there is only one address in the pool. If more than one address is configured by the user, the system does not raise any error, instead only the first address prefix of the pool is chosen for translation.
- For packets going from internal network to external network, if the internal subnet is not mapped or is set to 0xFFFF, then the datagram is discarded and an ICMP destination unreachable error is generated.
- For packets going from internal network to external network, if the 16-bit word has the adjustment added to it using the 1's complement method and is equal to 0xFFFF, then the value is written as zero.
- For packets coming from the external network to internal network, if the 16-bit word has the adjustment subtracted from it using 1's complement method and is equal to 0xFFFF, the 16-bit word is overwritten as zero.
- For translation of prefixes /48 or shorter, the adjustment must be added or subtracted from the first 16 bits after the /48 subnet mask, the values of which are not 0xFFFF. If the prefix is /49 or longer, then the adjustment must be added or subtracted from the first 16 bits (from 64 to 123), the values of which are not 0xFFFF.

**Related**   •  
**Documentation**

---

## Interoperation of Functionalities with Network Prefix Translation for IPv6

---

This topic contains the following sections that describe the working behavior of different functionalities with stateless source IPv6 prefix translation and the various system conditions:

### Address Mapping Algorithm

The NPTv6 translator filters the packets going out of the network and, if the source address of the packet matches with the source address defined in the rule (the **from** or source address in configuration), the source address is replaced with an address prefix from the pool defined for the rule. The next 16 bits after the prefix of the source address

are replaced with the checksum-adjusted value to ensure that the checksum remains the same in the outgoing packet even though the source address is changed. During the definition of the configuration rule and pool for the packets going outside the network, a denat rule and pool are created for the translation of the destination address for the packets coming into the network.

### Internal to External Translation

When a packet is going from the internal network to the external network, the IPv6 prefix in the source address of the packet (coming from inside node) is mapped to the external prefix. After checksum adjustment, the packet is routed toward the external network.

### External to Internal Translation

When a packet is coming from external network to internal network, the IPv6 prefix in the destination address of the packet (coming from outside host) is mapped to the internal prefix. After checksum adjustment, the packet is routed to internal network.

### Checksum-Neutral Translation

The NPTv6 translator translates the source address prefix in such a way that the transport layer checksum of the packet does not need to be recomputed. A checksum change caused by modifying part of the area covered by the checksum can be corrected by making an additional change to a different 16-bit field covered by the same checksum. This checksum neutral method first computes 1's complement checksum of the **internal-prefix** and the **external-prefix**.

For packets coming from the internal network, the adjustment is calculated as 1's complement and it is computed as follows:

Adjustment = Internal prefix checksum – External prefix checksum.

The adjustment value is added to the 16-bit word of the source address after the prefix.

For packets coming from external network, the adjustment is 1's complement and it is calculated as follows:

Adjustment = External prefix checksum – Internal prefix checksum.

The adjustment is added to the 16-bit word of the destination address after the translated prefix.

### Multihoming

If there are two NPTv6 translators with different external IPv6 prefix configurations for the same internal IPv6 prefix, then these two NPTv6 translators will translate the same internal IPv6 network prefix to two different external IPv6 network prefixes, depending on the translator the packet traverses.

### Hairpinning

When an internal node has knowledge of only the external (that is, the global address) of another internal node, it uses that address to send packet to that internal node. If such

a packet is received by an NPTv6 translator, that packet is routed toward the internal network again after it undergoes source address and destination address translation.

## Load Balancing

Load sharing is achieved when two translators have the same internal to external mapping configuration and packet load is shared between them. How the load balancing is achieved is beyond the scope of NPTv6.

The balancing could be implemented based on subnet ID portion of the IPv6 address. There can be two si- logical interfaces with the same mapping of the internal prefix to the external prefix. Packets are routed to one of the si- logical interfaces based on the subnet ID.

## ICMPv6 for NPTv6

Any host in the internal network should be able to ping a host in the external network through an NPTv6 translator. All ICMP error messages should be forwarded to the host in the internal network. During internal to external translation if there is no mapping possible for a prefix, then packet is dropped and the ICMP Destination Unreachable message is sent back. An ICMPv6 Destination Unreachable error is returned by the translator if the ICMP error is enabled in the following cases:

- If source address prefix lengths are less than or equal to /48 and the 48-63 bits are equal to 0xFFFF
- If prefix lengths are greater than /48 and all the 16-bit blocks in the interface ID (IID) (bits 64-127) are equal to 0xFFFF

Otherwise the packet is discarded.

For prefix lengths less than or equal to /48, the bits 48-63 are used as the 16-bit word adjusted for checksum. For prefix lengths greater than /48, the first 16-bit block in the IID that does not have the value 0xFFFF is the 16-bit word adjusted for checksum.

If the interface ID (IID) part of the address to be translated contains all zeros, ICMPv6 Parameter Problem error is returned by the translator if the ICMP error option is enabled. Otherwise the packet is discarded. ICMPv6 errors are generated by the control plane. The source address of the ICMPv6 packet is the IP address of the ingress interface.

### Related Documentation

- 

---

## Working of NPTv6 with Interface-Style and Next Hop-Style Service Sets

The objective is to add network prefix translation for IPv6 (NPTv6) inline service that performs stateless translation of the source IPv6 address. Consider a sample topology in which NPTv6 is implemented between an internal network with the prefix of FD01:0203:0405:/48 and an external network with the prefix of 2001:0DB8:0001:/48.

The source addresses FD01:0203:0405:/48 in the packets from a single administrative domain (internal network) destined to hosts in global network (external network) will

be translated to 2001:0DB8:0001::/48. Packets destined to internal network coming from external network will have their destination address as 2001:0DB8:0001::/48. This destination address will be mapped to internal network address FD01:0203:0405::/48 and will be forwarded to the internal network host. The lengths of both the subnets are assumed to be the same for this case. If they differ the shorter one would be extended to the prefix length of the longer one by suffixing zeros.

Address mapping algorithm used for NPTv6 is checksum-neutral. The translated IP headers will generate the same IPv6 pseudo-header checksum. Checksum is calculated using the standard Internet checksum algorithm. Changes that are made during translation of the IPv6 prefix are offset by the calculated changes made to the other parts of the IPv6 address. This results in transport layers that use the Internet checksum (such as TCP and UDP) calculating the same IPv6 pseudo-header checksum for both the internal and external forms of the same datagram and avoids the need to modify transport layer headers to correct the checksum value. The algorithm can map the addresses for inbound packets and outbound packets.

The NPTv6 translator works for both fragmented packets and packets with IP options enabled. The configuration changes required for NPTv6 are covered in the next sections.

The configuration of a router to handle services is through the definition of logical service interface, service sets and service set rules. These define how the service is applied to the packets.

The inline services logical interface, si-ifl, implementation available for static v4-v4 source-address inline-NAT can be reused for inline NPTv6. The configuration for the NPTv6 implemented for MS-DPC can be modified for inline NPTv6 implementation. There are two types of service set configurations—interface style and next hop style.

For the next hop-style service, a route entry is configured to steer packets to an inline service interface. There the packet would go through the service rules. If the packet matches the service rules, it is processed as per the service rules. For the interface-style service, the service set is configured directly on the media interface affecting traffic as it leaves and enters the interface. The packets are steered to the inline service interface by the service filter applied to the media interface.

#### Related Documentation

### Example: Achieving Address Independence By Configuring Stateless Network Prefix Translation in IPv6 Networks by Using Interface-Style Service Sets

You can configure stateless translation of source address prefixes in IPv6 networks (IPv6 to IPv6) on MX Series routers with MPCs that support inline NAT. The NPTv6 translator translates the source address prefix in such a way that the transport layer checksum of the packet does not need to be recomputed. NPTv6 defines a stateless method of IPv6 network prefix translation between internal and external networks. NPTv6 does not maintain the state for each node or each flow in the translator. You can use the **show services nat mappings nptv6 (internal | external)** command to view the NAT mappings for NPTv6 for internal and external addresses respectively. You can also use the **show**

**services inline nat statistics** and **show services inline nat pool** commands to display information about inline NAT with NPTv6 configured.



**NOTE:** This functionality is supported on MX Series routers with Trio-based FPCs (MPCs).

This example describes how to configure stateless source prefix translation for IPv6 packets using interface-style service sets on MX Series routers with MPCs, and contains the following sections:

- [Requirements on page 206](#)
- [Overview and Topology for Stateless Network Prefix Translation in IPv6 Networks Using Interface-Style Service Sets on page 206](#)
- [Configuration on page 207](#)
- [Verification on page 209](#)

## Requirements

This example uses the following hardware and software components:

- One MX Series router with an MPC.
- Junos OS Release 15.1R1 or later for MX Series routers

## Overview and Topology for Stateless Network Prefix Translation in IPv6 Networks Using Interface-Style Service Sets

For the interface style service, the service set is configured directly on the media interface affecting traffic as it leaves and enters the interface. The packets are steered to the inline service interface by the service filter applied to the media interface.

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.

Consider a sample configuration scenario in which NPTv6 is configured using interface-style service sets. An inline services interface, si-0/1/0, is configured with a bandwidth reserved for 10 gigabits per second. The si-0/1/0 interface is defined with inet6 family. A NAT address pool, nptv6\_pool, is configured with the address of abcd:ef12:3456::/48. A NAT rule is applied in the input direction to perform NPTv6 translation on packets that arrive from the source address of 1234:5678:9abc::/48. For packets from the source address of 1234:5678:9abc::/48 that match the NAT rule criterion, the address from the NAT address pool is allocated. A service set, ss\_nptv6, is specified with the NAT rule. A gigabit Ethernet interface, ge-5/0/0, is configured and the service set is applied to this interface.



## Configuration

To configure stateless network prefix translation for IPv6 using interface-style service sets, perform these tasks:

<b>CLI Quick Configuration</b>	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 <b>[edit]</b> hierarchy level:
<b>Configuring Interfaces</b>	<code>set interfaces si-0/1/0 unit 0 family inet6</code>
<b>Configuring Interfaces for Traffic to Be Handled By the Service Set</b>	<code>set interfaces ge-5/0/0 unit 0 family inet6 service input service-set nptv6-service-set</code> <code>set interfaces ge-5/0/0 unit 0 family inet6 service output service-set nptv6-service-set</code> <code>set interfaces ge-5/0/0 unit 0 family inet6 address 1234:5678:9abc::1/64</code>
<b>Configuring Bandwidth for the Service Inline (si-) Interface</b>	<code>set chassis fpc 0 pic 1 inline-services bandwidth 10g</code>
<b>Configuring NAT Pool and Rules</b>	<code>set services nat pool ss_nptv6_pool address abcd:ef12:3456::/48</code> <code>set services nat rule ss_nptv6_rule match-direction input term t0 from source-address 1234:5678:9abc::/48</code> <code>set services nat rule ss_nptv6_rule match-direction input term t0 then translated source-pool ss_nptv6_pool</code> <code>set services nat rule ss_nptv6_rule match-direction input term t0 then translated translation-type nptv6</code>
<b>Configuring the Service Set</b>	<code>set services service-set ss_nptv6 nat-rules ss_nptv6_rule</code> <code>set services service-set ss_nptv6 nat-options nptv6 icmpv6-error-messages</code> <code>set services service-set ss_nptv6 interface-service service-interface si-0/1/0.0</code>

**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 stateless network prefix translation for IPv6 using interface-style service sets:

1. Configure an inline services (si-) interface.  

```
[edit]
user@host# set interfaces si-0/1/0 unit 0 family inet6
```
2. Configure the interfaces for traffic to be handled by the service set.  

```
[edit]
user@host# set interfaces ge-5/0/0 unit 0 family inet6 service input service-set nptv6-service-set
user@host# set interfaces ge-5/0/0 unit 0 family inet6 service output service-set nptv6-service-set
user@host# set interfaces ge-5/0/0 unit 0 family inet6 address 1234:5678:9abc::1/64
```

3. Configure the bandwidth for the service inline (si-) interface.

```
[edit]
user@host# set chassis fpc 0 pic 1 inline-services bandwidth 10g
```

4. Configure a NAT pool and rule.

```
[edit]
user@host# set services nat pool ss_nptv6_pool address abcd:ef12:3456::/48
user@host# set services nat rule ss_nptv6_rule match-direction input term t0 from
source-address 1234:5678:9abc::/48
user@host# set services nat rule ss_nptv6_rule match-direction input term t0 then
translated source-pool ss_nptv6_pool
user@host# set services nat rule ss_nptv6_rule match-direction input term t0 then
translated translation-type nptv6
```

5. Configure the service set

```
[edit]
user@host# set services service-set ss_nptv6 nat-rules ss_nptv6_rule
user@host# set services service-set ss_nptv6 nat-options nptv6
icmpv6-error-messages
user@host# set services service-set ss_nptv6 interface-service service-interface
si-0/1/0.0
```

---

## Results

From the configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces**, and **show services** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show chassis
chassis {
  fpc 0 {
    pic 1 {
      inline-services {
        bandwidth 10g;
      }
    }
  }
}
```

```
user@host# show interfaces
interfaces {
  si-0/1/0 {
    unit 0 {
      family inet6;
    }
  }
  ge-5/0/0 {
    unit 0 {
      family inet6 {
        service {
          input {
            service-set nptv6-service-set;
          }
          output {
            service-set nptv6-service-set;
          }
        }
      }
    }
  }
}
```

```

    }
address 1234:5678:9abc::1/64;
}
}
}

user@host# show services
services {
  service-set ss_nptv6 {
    nat-rules ss_nptv6_rule;
    nat-options {
      nptv6 {
        icmpv6-error-messages;
      }
    }
    interface-service {
      service-interface si-0/1/0.0;
    }
  }

  nat {
    pool ss_nptv6_pool {
      address abcd:ef12:3456::/48;
    }
    rule ss_nptv6_rule {
      match-direction input;
      term t0 {
        from {
          source-address {
            1234:5678:9abc::/48;
          }
        }
        then {
          translated {
            source-pool ss_nptv6_pool;
            translation-type {
              nptv6;
            }
          }
        }
      }
    }
  }
}

```

## Verification

To confirm that the configuration is working properly, perform the following:

- [Verifying the NAT Pool Mappings on page 209](#)
- [Verifying the Inline NAT Pools and Statistics on page 210](#)

### Verifying the NAT Pool Mappings

**Purpose** Verify the existing NAT address pools and mappings for IPv6 network prefix translation.

**Action** From operational mode, use the **show services nat mappings nptv6** command:

```
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

```
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

**Meaning** The output shows the mapping between NAT addresses and ports for IPv6 stateless network prefix translation of external and internal addresses. The address and port details that are originally sent and converted using NAT are displayed.

### Verifying the Inline NAT Pools and Statistics

**Purpose** Verify the inline NAT pools and statistics for IPv6 network prefix translation.

**Action** From operational mode, use the **show services inline nat** command:

```
user@host> show services inline nat statistics interface si-4/0/0
```

```
Service PIC Name
:si-4/0/0

Control Plane Statistics
ICMPv4 errors packets pass through          :0
ICMPv4 errors packets locally generated      :0
ICMPv6 errors packets pass through          :0
ICMPv6 errors packets locally generated      :0
Dropped packets                             :0

Data Plane Statistics
NATed packets                               :0
deNATed packets                             :0
Errors                                       :0
```

```
user@host> show services inline nat pool
```

```
Interface: si-4/0/0, Service set: ss_nptv6
NAT pool: ss_nptv6_pool1, Translation type: NPTV6
Address range: abcd:ef12:3456::/48
NATed packets: 0, deNATed packets: 0, Errors: 0

NAT pool: ss_nptv6_pool2, Translation type: NPTV6
Address range: 1111:2222:3333::/48
NATed packets: 0, deNATed packets: 0, Errors: 0
```

```
user@host> show services inline nat pool ss_nptv6_pool1
```

```

Interface: si-4/0/0, Service set: ss_nptv6
NAT pool: ss_nptv6_pool1, Translation type: NPTv6
Address range: abcd:ef12:3456::/48
NATed packets: 0, deNATed packets: 0, Errors: 0

```

**Meaning** The output shows the information about inline NAT address translations, such as the number of packets that are subject to NAT processing, the packets that are not translated, and the packets with translation errors for a specified service set and an si- interface.

**Related** •  
**Documentation**

## Example: Achieving Address Independence By Configuring Stateless Network Prefix Translation in IPv6 Networks by Using Next-Hop -Style Service Sets

You can configure stateless translation of source address prefixes in IPv6 networks (IPv6 to IPv6) on MX Series routers with MPCs where inline NAT is supported. The NPTv6 translator translates the source address prefix in such a way that the transport layer checksum of the packet does not need to be recomputed. NPTv6 defines a stateless method of IPv6 network prefix translation between internal and external networks. NPTv6 does not maintain per node or per flow state in the translator. You can use the `show services nat mappings nptv6 (internal | external)` command to view the NAT mappings for NPTv6 for internal and external addresses respectively. You can also use the `show services inline nat statistics` and `show services inline nat pool` commands to display information about inline NAT with NPTv6 configured.



**NOTE:** This functionality is supported on MX Series routers with Trio-based FPCs (MPCs).

This example describes how to configure stateless source prefix translation for IPv6 packets using next-hop style service sets on MX Series routers with MPCs, and contains the following sections:

- [Requirements on page 211](#)
- [Overview and Topology of Stateless Network Prefix Translation in IPv6 Networks Using Next-Hop Style Service Sets on page 212](#)
- [Configuration on page 212](#)
- [Verification on page 216](#)

### Requirements

This example uses the following hardware and software components:

- One MX Series router with an MPC.
- Junos OS Release 15.1R1 or later for MX Series routers

## Overview and Topology of Stateless Network Prefix Translation in IPv6 Networks Using Next-Hop Style 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.

For the next hop style service, a route entry is configured to steer packets to an inline service interface. The packet is validated through the service rules. If the packet matches the service rules, it would be processed according to the service rules.

Consider a sample configuration scenario in which NPTv6 is configured using next-hop style service sets. An inline services interface, si-0/1/0, is configured with a bandwidth reserved for 10 gigabits per second. The si-0/1/0 interface is defined with inet6 family. A NAT address pool, nptv6\_pool, is configured with the address of abcd:ef12:3456::/48. A NAT rule is applied in the input direction to perform NPTv6 translation on packets that arrive from the source address of 1234:5678:9abc::/48. For packets from the source address of 1234:5678:9abc::/48 that match the NAT rule criterion, the address from the NAT address pool is allocated. The service set is configured for forwarding next-hops with the service interface of si-0/1/0.1 associated with the service set applied inside the network, with parameters for next hop service interfaces for the inside network and si-/1/0.2 associated with the service set applied outside the network. A service set, ss\_nptv6, is specified with the NAT rule. The service interface domain is specified for the si- interface with the inside service-domain configured for si-0/1/0.1 and outside service domain configured for si-0/1/0.2. A routing instance, inst1, is configured with the instance type as a VRF instance. interface si-0/1/0.1 and interface ge-5/0/0 are associated with inst1. The inside and outside interface domain matches that specified with the inside-service-interface and outside-service-interface statements. A policy is configured for NAT events with the action to reject all packets.

## Configuration

To configure stateless network prefix translation for IPv6 using next-hop style service sets, perform these tasks:

<b>CLI Quick Configuration</b>	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 <b>[edit]</b> hierarchy level:
<b>Configuring Inline Interfaces</b>	<pre>set interfaces si-0/1/0 unit 0 family inet6 set interfaces si-0/1/0 unit 1 family inet6 set interfaces si-0/1/0 unit 1 service-domain inside set interfaces si-0/1/0 unit 2 family inet6 set interfaces si-0/1/0 unit 2 service-domain outside set interfaces ge-5/0/0 unit 0 family inet6 address 1234:5678:9abc::1/64</pre>
<b>Configuring Bandwidth for Inline Services</b>	<pre>set chassis fpc 0 pic 1 inline-services bandwidth 10g</pre>

Configuring NAT Pool and Rule	<pre> set services nat pool ss_nptv6_pool address abcd:ef12:3456::/48 set services nat rule ss_nptv6_rule match-direction input term t0 from source-address 1234:5678:9abc::/48 set services nat rule ss_nptv6_rule match-direction input term t0 then translated source-pool ss_nptv6_pool set services nat rule ss_nptv6_rule match-direction input term t0 then translated translation-type nptv6 </pre>
Configuring a Service Set	<pre> set services service-set ss_nptv6 nat-rules ss_nptv6_rule set services service-set ss_nptv6 nat-options nptv6 icmpv6-error-messages set services service-set ss_nptv6 nexthop-service inside-service-interface si-0/1/0.1 set services service-set ss_nptv6 nexthop-service outside-service-interface si-0/1/0.2 </pre>
Configuring Routing Instances	<pre> set routing-instances inst1 instance-type vrf set routing-instances inst1 interface si-0/1/0.1 set routing-instances inst1 interface ge-5/0/0.0 set routing-instances inst1 route-distinguisher 1234:5678 set routing-instances inst1 vrf-import reject-all set routing-instances inst1 vrf-export reject-all set routing-instances inst1 routing-options rib inst1.inet6.0 static route ::0/0 next-hop si-0/1/0.1 </pre>
Configuring the Policy and Action Modifier	<pre> set policy-options policy-statement reject-all then reject </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> <p>To configure stateless network prefix translation for IPv6 using next-hop style service sets:</p> <ol style="list-style-type: none"> <li>1. Configure the inline interface for NAT services. <pre> [edit] user@host# set interfaces si-0/1/0 unit 0 family inet6 user@host# set interfaces si-0/1/0 unit 1 family inet6 user@host# set interfaces si-0/1/0 unit 1 service-domain inside user@host# set interfaces si-0/1/0 unit 2 family inet6 user@host# set interfaces si-0/1/0 unit 2 service-domain outside user@host# set interfaces ge-5/0/0 unit 0 family inet6 address 1234:5678:9abc::1/64 </pre> </li> <li>2. Set the bandwidth for inline services. <pre> [edit] user@host# set chassis fpc 0 pic 1 inline-services bandwidth 10g </pre> </li> <li>3. Configure the NAT pool and rule. <pre> [edit] user@host# set services nat pool ss_nptv6_pool address abcd:ef12:3456::/48 user@host# set services nat rule ss_nptv6_rule match-direction input term t0 from source-address 1234:5678:9abc::/48 user@host# set services nat rule ss_nptv6_rule match-direction input term t0 then translated source-pool ss_nptv6_pool </pre> </li> </ol>

```
user@host# set services nat rule ss_nptv6_rule match-direction input term t0 then
translated translation-type nptv6
```

4. Configure a service set using the NAT rule associated with the NAT pool.

```
[edit]
user@host# set services service-set ss_nptv6 nat-rules ss_nptv6_rule
user@host# set services service-set ss_nptv6 nat-options nptv6
icmpv6-error-messages
user@host# set services service-set ss_nptv6 nexthop-service
inside-service-interface si-0/1/0.1
user@host# set services service-set ss_nptv6 nexthop-service
outside-service-interface si-0/1/0.2
```

5. Configure routing instances that use the si- interfaces configured.

```
[edit]
user@host# set routing-instances inst1 instance-type vrf
user@host# set routing-instances inst1 interface si-0/1/0.1
user@host# set routing-instances inst1 interface ge-5/0/0.0
user@host# set routing-instances inst1 route-distinguisher 1234:5678
user@host# set routing-instances inst1 vrf-import reject-all
user@host# set routing-instances inst1 vrf-export reject-all
user@host# set routing-instances inst1 routing-options rib inst1.inet6.0 static route
::0/0 next-hop si-0/1/0.1
```

6. Configure the policy and the action modifier for NAT packets.

```
[edit]
user@host# set policy-options policy-statement reject-all then reject
```

## Results

From the configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces**, **show policy-options**, **show routing-instances**, and **show services** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show chassis
chassis {
  fpc 0 {
    pic 1 {
      inline-services {
        bandwidth 10g;
      }
    }
  }
}

user@host# show interfaces

chassis {
  fpc 0 {
    pic 1 {
      inline-services {
        bandwidth 10g;
      }
    }
  }
}
```



```

interfaces {
  si-0/1/0 {
    unit 0 {
      family inet6;
    }
    unit 1 {
      family inet6;
      service-domain inside;
    }
    unit 2 {
      family inet6;
      service-domain outside;
    }
  }
  ge-5/0/0 {
    unit 0 {
      family inet6 {
        address 1234:5678:9abc::1/64;
      }
    }
  }
}

```

```

user@host# show policy-options
policy-options {
  policy-statement reject-all {
    then reject;
  }
}

```

```

user@host# show routing-instances
routing-instances {
  inst1 {
    instance-type vrf;
    interface si-0/1/0.1;
    interface ge-5/0/0.0;
    route-distinguisher 1234:5678;
    vrf-import reject-all;
    vrf-export reject-all;
    routing-options {
      rib inst1.inet6.0 {
        static {
          route ::0/0 next-hop si-0/1/0.1;
        }
      }
    }
  }
}

```

```

user@host# show services
services {
  service-set ss_nptv6 {
    nat-rules ss_nptv6_rule;
    nat-options {
      nptv6 {
        icmpv6-error-messages;
      }
    }
  }
}

```

```

}
nexthop-service {
  inside-service-interface si-0/1/0.1;
  outside-service-interface si-0/1/0.2;
}
}
nat {
  pool ss_nptv6_pool {
    address abcd:ef12:3456::/48;
  }
  rule ss_nptv6_rule {
    match-direction input;
    term t0 {
      from {
        source-address {
          1234:5678:9abc::/48;
        }
      }
      then {
        translated {
          source-pool ss_nptv6_pool;
          translation-type {
            nptv6;
          }
        }
      }
    }
  }
}
}
}
}
}
}
}

```

## Verification

To confirm that the configuration is working properly, perform the following:

- [Verifying the NAT Pool Mappings on page 216](#)
- [Verifying the Inline NAT Pools and Statistics on page 217](#)

### Verifying the NAT Pool Mappings

**Purpose** Verify the existing NAT address pools and mappings for IPv6 network prefix translation.

**Action** From operational mode, use the **show services nat mappings nptv6** command:

```
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

```
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

**Meaning** The output shows the information about inline NAT address translations, such as the number of packets that are subject to NAT processing, the packets that are not translated, and the packets with translation errors for a specified service set and an si- interface.

### Verifying the Inline NAT Pools and Statistics

**Purpose** Verify the inline NAT pools and statistics for IPv6 network prefix translation.

**Action** From operational mode, use the **show services inline nat** command:

```
user@host> show services inline nat statistics interface si-4/0/0
```

```
Service PIC Name
      :si-4/0/0
```

#### Control Plane Statistics

```
ICMPv4 errors packets pass through      :0
ICMPv4 errors packets locally generated :0
ICMPv6 errors packets pass through      :0
ICMPv6 errors packets locally generated :0
Dropped packets                          :0
```

#### Data Plane Statistics

```
NATed packets
      :0
deNATed packets
      :0
Errors
      :0
```

```
user@host> show services inline nat pool
```

```
Interface: si-0/1/0, Service set: ss_nptv6
NAT pool: ss_nptv6_pool1, Translation type: NPTV6
Address range: abcd:ef12:3456::/48
NATed packets: 0, deNATed packets: 0, Errors: 0
```

```
NAT pool: ss_nptv6_pool2, Translation type: NPTV6
Address range: 1111:2222:3333::/48
NATed packets: 0, deNATed packets: 0, Errors: 0
```

```
user@host> show services inline nat pool ss_nptv6_pool1
Interface: si-0/1/0, Service set: ss_nptv6
NAT pool: ss_nptv6_pool1, Translation type: NPTV6
Address range: abcd:ef12:3456::/48
NATed packets: 0, deNATed packets: 0, Errors: 0
```

**Meaning** The output shows the mapping between NAT addresses and ports for IPv6 stateless network prefix translation of external and internal addresses. The address and port details that are originally sent and converted using NAT are displayed.

**Related Documentation** •



# Monitoring NAT

- [Configuring NAT Session Logs on page 219](#)
- [Monitoring NAT Pool Usage on page 220](#)

## 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 26](#)
  - [Interim Logging for Port Block Allocation on page 163](#)

---

## 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 3

# Transitioning to IPv6 Using Softwires

- [Softwires Overview on page 223](#)
- [Softwires Configuration Overview on page 229](#)
- [Transitioning to IPv6 Using 6to4 Softwires on page 233](#)
- [Transitioning to IPv6 Using DS-Lite Softwires on page 237](#)
- [Transitioning to IPv6 Using 6rd Softwires on page 255](#)
- [Monitoring and Troubleshooting Softwires on page 293](#)





# Softwires Overview

- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 223](#)

## Tunneling Services for IPv4-to-IPv6 Transition Overview

---

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 requires 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 the packet 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. The number of established softwires does not affect throughput, and scalability is independent of the number of interfaces. Scalability is only limited to the number of flows that the services DPC or PIC can support.

This topic contains the following sections:

- [6to4 Overview on page 223](#)
- [DS-Lite Softwires—IPv4 over IPv6 on page 225](#)
- [6rd Softwires—IPv6 over IPv4 on page 226](#)

### 6to4 Overview

- [Basic 6to4 on page 224](#)
- [6to4 Anycast on page 224](#)
- [6to4 Provider-Managed Tunnels on page 225](#)

## Basic 6to4

---

6to4 is an Internet transition mechanism for migrating from IPv4 to IPv6, a system that enables 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, because IPv6 is not required on nodes between the host and the destination. 6to4 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, 6to4 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, the host's IPv6 default gateway must be set to a 6to4 address that 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. When processed by 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. We recommend that providers who want to provide 6to4 service to their clients or peers advertise the Anycast prefix like any other IP prefix, and route the prefix to the provider's 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 the embedding of IPv4 routes into the routing tables of IPv6 routers. From the 6to4 relay router the packets can then be sent over the IPv4 Internet to the destination.

## 6to4 Anycast

---

Router 6to4 requires 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. Removing this configuration makes the solution available for small or domestic users, even those with a single host or simple home gateway instead of a border router. Removing this configuration 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 providers can implement 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 *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 that translates the source 6to4 prefix to a provider assigned prefix that 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 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 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, 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 *softwire initiator* has been replaced by *B4*.
- The term *softwire 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 softwire flow is shown in [Figure 17 on page 226](#).

**Figure 17: 6rd Softwire Flow**



Junos OS supports a 6rd softwire concentrator on a services DPC or PIC to facilitate rapid deployment of IPv6 service to subscribers on native IPv4 customer edge WANs. IPv6 packets are encapsulated in IPv4 packets by a softwire initiator at the customer edge WAN. These packets are tunneled to a softwire concentrator residing on an MS-DPC (branch relay). A softwire is created when IPv4 packets containing IPv6 destination information are received at the softwire 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 softwire and sent to the customer edge WAN.

The softwire concentrator creates softwires as the IPv4 packets are received from the customer edge WAN side or IPv6 packets are received from the Internet. A 6rd softwire on the Services DPC is identified by the 3-tuple containing the service set ID, customer edge softwire initiator IPv4 address, and softwire concentrator IPv4 address. IPv6 flows

are also created for the encapsulated IPv6 payload, and are associated with the specific softwire 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 routers equipped with Multiservices DPCs.

Junos OS supports inline 6rd and 6to4 on Modular Port Concentrator (MPC) line cards.

For more information on 6rd softwires, see RFC 5969, *IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification*.

**Related  
Documentation**

- [Junos Address Aware Network Addressing Overview on page 35](#)
- [Configuring a 6rd Softwire Concentrator on page 255](#)
- [Configuring a DS-Lite Softwire Concentrator on page 237](#)
- [Configuring Softwire Rules on page 229](#)
- [Configuring Service Sets for Softwire on page 230](#)
- [Configuring Inline 6rd on page 279](#)



# Softwires Configuration Overview

- [Configuring Software Rules on page 229](#)
- [Configuring Service Sets for Software on page 230](#)

## 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 223](#)
- [Configuring a 6rd Software Concentrator on page 255](#)
- [Configuring a DS-Lite Software Concentrator on page 237](#)
- [Configuring IPv6 Multicast Interfaces on page 238](#)
- [Configuring Service Sets for Software on page 230](#)

---

## 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 softwires.



---

**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 14.”



- Related Documentation**
- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 223](#)
  - [Configuring Softwire Rules on page 229](#)
  - [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 244](#)



# Transitioning to IPv6 Using 6to4 Softwires

- [Configuring a 6to4 Provider-Managed Tunnel on page 233](#)

## 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 237](#)
- [Configuring IPv6 Multicast Interfaces on page 238](#)
- [Example: Basic DS-Lite Configuration on page 238](#)
- [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 244](#)
- [Protecting CGN Devices Against Denial of Service \(DOS\) Attacks on page 251](#)
- [DS-Lite Subnet Limitation on page 252](#)

## 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 software

```
[edit services software 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 223](#)
- [Configuring Software Rules on page 229](#)
- [Configuring IPv6 Multicast Interfaces on page 238](#)
- [Configuring Service Sets for Software on page 230](#)
- [Example: Basic DS-Lite Configuration on page 238](#)
- [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 244](#)

---

## Configuring IPv6 Multicast Interfaces

---

Configure multicast filters on Ethernet interfaces when IPv6 NAT is used for neighbor discovery. This enables the router to process software-initiated flows in both directions.

To configure IPv6 multicast interfaces:

1. Access the software hierarchy.

```
user@host# edit services software
```

2. Include the [ipv6-multicast-interfaces](#) statement for an individual interface.

```
[edit services software]  
user@host# set ipv6-multicast-interfaces interface-name
```

Or configure all software interfaces as IPv6 multicast.

```
[edit services software]  
user@host# set ipv6-multicast-interfaces all
```

**Related  
Documentation**

- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 223](#)
- [Configuring a 6rd Software Concentrator on page 255](#)
- [Configuring a DS-Lite Software Concentrator on page 237](#)
- [Configuring Software Rules on page 229](#)

---

## Example: Basic DS-Lite Configuration

---

- [Requirements on page 239](#)
- [Configuration Overview and Topology on page 239](#)
- [Configuration on page 239](#)



## 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 18 on page 239](#).

**Figure 18: 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 240](#)
- [Interfaces Configuration on page 240](#)
- [Network Address and Port Translation Configuration on page 241](#)
- [Softwire Configuration on page 242](#)
- [Service Set Configuration on page 243](#)

### 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 (software initiator) and facing the Internet:

1. Go to 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:

- Go to the **[edit services nat]** hierarchy level.  

```

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

```
- Define a NAT pool p1.  

```

user@host# set pool p1 address 129.0.0.1/32 port automatic

```
- Define a NAT rule, beginning with the match direction.  

```

[edit services nat]
user@host# set rule r1 match-direction input

```
- 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

```
- 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

```
- (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 223](#)
  - [Configuring a DS-Lite Software Concentrator on page 237](#)
  - [Configuring Software Rules on page 229](#)
  - [Configuring Service Sets for Software on page 230](#)
  - [Example: Basic 6rd Configuration on page 257](#)
  - [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 244](#)

---

## Example: Configuring DS-Lite and 6rd in the Same Service Set

- [Requirements on page 244](#)
- [Overview on page 244](#)
- [Configuration on page 244](#)

### 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@host# 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@host# 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@host# 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@host# 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 223](#)
  - [Configuring Service Sets for Softwire on page 230](#)
  - [Example: Basic DS-Lite Configuration on page 238](#)
  - [Example: Basic 6rd Configuration on page 257](#)

## 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 251](#)
- [EIF Inbound Flow Limit on page 252](#)

### 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 252](#)
- [Configuring DS-Lite Per Subnet Session Limitation to Prevent Denial of Service Attacks on page 253](#)

### 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 software 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@host# set services service-set service-set-name software-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@host# set services software software-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 255](#)
- [Configuring Stateful Firewall Rules for 6rd Software on page 256](#)
- [Example: Basic 6rd Configuration on page 257](#)
- [Inter-Chassis High Availability for MS-MIC and MS-MPC on page 262](#)
- [High Availability and Load Balancing for 6rd Softwires on page 274](#)
- [Configuring Inline 6rd on page 279](#)
- [Inline 6rd and 6to4 Configuration Guidelines on page 283](#)
- [Examples: 6rd and 6to4 Configurations on page 284](#)

## 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 223](#)
- [Configuring Software Rules on page 229](#)
- [Configuring Stateful Firewall Rules for 6rd Software on page 256](#)
- [Configuring Service Sets for Software on page 230](#)
- [Example: Basic 6rd Configuration on page 257](#)
- [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 244](#)

---

## 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 223](#)
- [Configuring a 6rd Software Concentrator on page 255](#)
- [Configuring Software Rules on page 229](#)
- [Configuring Service Sets for Software on page 230](#)
- [Example: Basic 6rd Configuration on page 257](#)
- [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 244](#)

## Example: Basic 6rd Configuration

- [Requirements on page 257](#)
- [Overview on page 257](#)
- [Configuration on page 257](#)

### 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@host# 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@host# **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 223](#)
  - [Configuring a 6rd Software Concentrator on page 255](#)
  - [Configuring Software Rules on page 229](#)
  - [Configuring Stateful Firewall Rules for 6rd Software on page 256](#)
  - [Configuring Service Sets for Software on page 230](#)
  - [Example: Basic DS-Lite Configuration on page 238](#)
  - [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 244](#)

## 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 262](#)
- [Configuring Inter-Chassis High Availability for Stateful Firewall and NAPT44 \(MS-MPC, MS-MIC\) on page 263](#)
- [Example: Inter-Chassis Stateful High Availability for NAT and Stateful Firewall \(MS-MIC, MS-MPC\) on page 264](#)

### 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 19 on page 263](#) shows the inter-chassis high availability topology.



Figure 19: 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-name 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 264](#)
- [Overview on page 264](#)
- [Configuration on page 265](#)

### 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

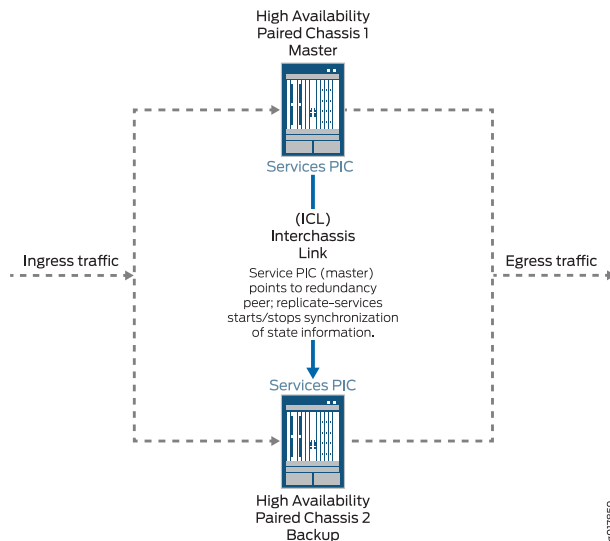
---

Two MX 3D routers are identically configured to facilitate stateful failover for firewall and NAT services in case of a chassis failure.

### Topology

[Figure 20 on page 265](#) shows the inter-chassis high availability topology.

Figure 20: Inter-Chassis High Availability Topology



### Configuration

To configure inter-chassis high availability for this example, perform these tasks:

- [Configuring Interfaces for Chassis 1 on page 267](#)
- [Configure Routing Information for Chassis 1 on page 268](#)
- [Configuring NAT and Stateful Firewall for Chassis 1 on page 269](#)
- [Configuring the Service Set on page 270](#)
- [Configuring Interfaces for Chassis 2 on page 271](#)
- [Configure Routing Information for Chassis 2 on page 273](#)

### 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 5.5.5.1/32 next-hop
  ms-4/0/0.10
set routing-instances HA routing-options static route route 5.5.5.2/32 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 5.5.5.2/32
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 274](#)
- [Example: Load Balancing a 6rd Domain Across Multiple Services PICs on page 274](#)
- [Configuring High Availability for 6rd Using 6rd Anycast on page 279](#)

### 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 274](#)
- [Overview on page 275](#)
- [Configuration on page 275](#)

---

#### 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 275](#)
- [Software Concentrator and Software Rule Configuration on page 276](#)
- [Stateful Firewall Configuration on page 276](#)
- [Service Set Configuration on page 277](#)
- [Load-Balancing Configuration on page 277](#)

### 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
```

### *Softwire Concentrator and Softwire Rule Configuration*

**Step-by-Step Procedure** The softwire configuration is straightforward. In this example, the 6rd domain prefix is 3040::0/16, the 6rd softwire 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 softwire:

1. Go to the **[edit services softwire]** hierarchy level.  
`user@host# edit services softwire`
2. Configure IPv6 multicast.  
`[edit services softwire]  
user@host# set ipv6-multicast-interfaces all`
3. Go to the softwire concentrator v6rd hierarchy level and name the softwire concentrator **shenick01-rd1**.  
`[edit services softwire]  
user@host# edit softwire-concentrator v6rd shenick01-rd1`
4. Configure the softwire concentrator properties.  
`[edit services softwire softwire-concentrator v6rdshenick01-rd1 ]  
user@host# set softwire-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 softwire rule for incoming 6rd traffic.  
`[edit services softwire softwire-concentrator v6rd shenick01-rd1 ]  
user@host# up 1  
[edit services softwire ]  
user@host# edit rule shenick01-r1  
[edit services softwire 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@host# 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@host# 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*

## Configuring Inline 6rd

Junos OS supports inline 6rd and 6to4 on Modular Port Concentrator (MPC) line cards, saving customers the cost of using Multiservices Dense Port Concentrators (MS-DPCs) for the required tunneling, encapsulation, and decapsulation processes. Anycast is supported for 6to4 (next-hop service interfaces only). Hairpinning is also supported for traffic between 6rd domains.

To implement the inline functionality, you configure service interfaces on the MPC as inline services interfaces (si-) rather than as multiServices (ms-) interfaces.

- [Configuring the Bandwidth for Inline Services on page 280](#)
- [Configuring the Interfaces on page 280](#)
- [Configuring the Software Concentrator and Rule on page 282](#)
- [Configuring the Service Set on page 282](#)
- [Configuring the Routing Instance on page 283](#)

## Configuring the Bandwidth for Inline Services

You must provide bandwidth configuration for inline services on the modular port concentrator (MPC) used for inline 6rd processing.

To configure bandwidth:

- Specify the MPC and logical interface, and the desired bandwidth, 1g or 10g.  

```
user@host# set chassis fpc mpc-number pic logical-interface-number inline-services  
bandwidth bandwidth
```

For example:

```
user@host# set chassis fpc 0 pic 0 inline-services bandwidth 10g
```

## Configuring the Interfaces

Configure the si- interfaces for 6rd control and data. 6rd services must be configured on port 0.

To configure the si- interfaces:

1. Configure the 6rd services on port 0 and include units for IPv4 and IPv6.

```
user@host# set interfaces si-mpc-number/logical-interface-number/0 unit 0 family  
inet  
user@host# set interfaces si-mpc-number/logical-interface-number/0 unit 0 family  
inet6
```

For example:

```
user@host# set interfaces si-0/0/0 unit 0 family inet  
user@host# set interfaces si-0/0/0 unit 0 family inet6
```

2. Configure the media interfaces for the inside service domain.

```
user@host# set interfaces si-mpc-number/logical-interface-number/0 unit unit-number  
family inet  
user@host# set interfaces si-mpc-number/logical-interface-number/0 unit unit-number  
family inet6  
user@host# set interfaces si-0/0/0 unit unit-number service-domain inside
```

For example:

```
user@host# set interfaces si-0/0/0 unit 1 family inet  
user@host# set interfaces si-0/0/0 unit 1 family inet6  
user@host# set interfaces si-0/0/0 unit 1 service-domain inside
```

3. Configure the media interfaces for the outside service domain.

```

user@host# set interfaces si-mpc-number/logical-interface-number/0 unit unit-number
family inet
user@host# set interfaces si-mpc-number/logical-interface-number/0 unit unit-number
family inet6
user@host# set interfaces si-mpc-number/logical-interface-number/0 unit unit-number
service-domain outside

```

For example:

```

user@host# set interfaces si-0/0/0 unit 2 family inet
user@host# set interfaces si-0/0/0 unit 2 family inet family inet6
user@host# set interfaces si-mpc-number/logical-interface-number/0 unit unit-number
service-domain outside

```

4. Configure the IPv4-facing interface for use with an interface-style or next-hop service set.

- To configure for use with an interface-style service set, configure input and output service and specify the service set.

```

user@host# set interfaces ge-mpc-number/logical-interface-number/port unit
unit-number family inet service input service-set service-set-name
user@host# set interfaces ge-mpc-number/logical-interface-number/port unit
unit-number family inet service output service-set service-set-name
user@host# set interfaces ge-mpc-number/logical-interface-number/port unit
unit-number family inet address ip-address

```

For example:

```

user@host# set interfaces ge-0/2/7 unit 0 family inet service input service-set
vrf-intf-service-set
user@host# set interfaces ge-0/2/7 unit 0 family inet service output service-set
vrf-intf-service-set
user@host# set interfaces ge-0/2/7 unit 0 family inet address 10.10.10.1/16

```

- To configure for use with a next-hop style service set, omit the **service input** and **service output** references.

```

user@host# set interfaces ge-mpc-number/logical-interface-number/port unit
unit-number family inet
user@host# set interfaces ge-mpc-number/logical-interface-number/port unit
unit-number family inet address ip-address

```

For example:

```

user@host# set interfaces ge-0/2/7 unit 0 family inet
user@host# set interfaces ge-0/2/7 unit 0 family inet address 10.10.10.1/16

```

5. Configure the IPv6 facing interface.

```

user@host# set interfaces ge-mpc-number/logical-interface-number/port unit
unit-number family inet6 address ipv6-address

```

For example:

```

user@host# set interfaces ge-0/2/8 unit 0 family inet6 address 3abc::1/16

```

## Configuring the Software Concentrator and Rule

Define the software concentrator and rule used for encapsulation and decapsulation of IPv6 over IPv4 packets for CE.

To define the software concentrator:

1. Specify a 6rd software concentrator and its address.

```
user@host# set services software software-concentrator v6rd concentrator-name
software-address ip-address
```

For example:

```
user@host# set services software software-concentrator v6rd swire01-rd1
software-address 30.30.30.1
```

2. Configure the IPv4 address prefix for the customer edge network and the IPv6 address prefix for the 6rd domain.

```
user@host# set services software software-concentrator v6rd concentrator-name
ipv4-prefix ipv4-prefix v6rd-prefix v6rd-prefix
```

For example:

```
user@host# set services software software-concentrator v6rd swire01-rd1 ipv4-prefix
10.10.0.0/16 v6rd-prefix 3040::0/16
```

3. Configure the size, in bytes, of the maximum transmission unit **mtu-ipv4** for IPv6 packets encapsulated in IPv4. Compute this as the maximum expected IPv4 packet size plus 20.

```
user@host# set services software software-concentrator v6rd concentrator-name set
mtu-ipv4 number-of-bytes
```

For example:

```
user@host# set services software software-concentrator v6rd swire01-rd1 set mtu-ipv4
9192
```

To configure the software rule:

- Specify the software rule, specifying the direction of traffic to be tunneled and the 6rd software concentrator to be used.

```
user@host# set services software rule software-rule-name match-direction
match-direction term rule-term-number then v6rd concentrator-name
```

For example:

```
user@host# set services software rule swire01-r1 match-direction input term t1 then
v6rd swire01-rd1
```

## Configuring the Service Set

To configure an interface style or next-hop service set for 6rd processing:

- Specify an interface style service set.

```
user@host# set services service-set service-set-name software-rules software-rule-name
service-interface interface-name
```

For example:

```
user@host# set services service-set vrf-intf-service-set software-rules swire01-r1
service-interface si-0/0/0.0
```

or

- Configure a next-hop service set.

```
user@host# set services service-set service-set-name software-rules software-rule-name
user@host# set services service-set service-set-name next-hop-service
inside-service-interface inside-interface outside-service-interface outside-interface

user@host# set services service-set vrf-nh-service-set software-rules swire01-r1
user@host# set services service-set vrf-nh-service-set next-hop-service
inside-service-interface si-0/0/0.1 outside-service-interface si-0/0/0.2
```

## Configuring the Routing Instance

To configure the routing instance:

1. Specify the routing instance and each interface it serves.

```
user@host# set routing-instance routing-instance-name instance-type vrf interface
interface-name
```

For example:

```
user@host# set routing-instance v6rd-vrf instance-type vrf interface si-0/0/0.1
user@host# set routing-instance v6rd-vrf instance-type vrf interface interface
ge-0/2/7.0
```

2. Specify the route distinguisher and vrf-target.

```
user@host# set routing-instance v6rd-vrf route-distinguisher 1.1.1.1:1
user@host# set routing-instance v6rd-vrf vrf-target target:100:100
```

### Related Documentation

- [Configuring a 6rd Software Concentrator on page 255](#)
- [Configuring Software Rules](#)
- [Configuring Inline 6rd on page 279](#)

## Inline 6rd and 6to4 Configuration Guidelines

Keep the following points in mind when you are configuring and using inline 6rd and 6to4:

- You can configure a maximum of 1024 software concentrators on a single line card.
- Reassembly of 6rd IPv4 packet from CE is not added as part of this release.
- 6rd multicast is not supported.
- Any ICMPv4 errors generated in the IPv4 access network (between CPE and border relays) are dropped on the border relay. They are not converted into IPv6 errors and forwarded to IPv6 side.

- 6rd/6to4 Anycast and load balancing can be configured only using next-hop style service-interface configuration, not interface style.
- The si- interface input features are not exercised for packets flowing to the 6rd tunnel.
- Bandwidth for traffic from the 6rd tunnel bandwidth is limited by the available PFE bandwidth; bandwidth for traffic to the 6rd tunnel is limited by the configured SI-IFD loopback bandwidth.
- If the packet length is more than Tunnel MTU for downlink packets after encapsulating with an IPv4 header, the packet is dropped as v4 MTU errors. For these packet drops an **ICMPv6 packet too big error** message is sent back to the sender. Typically 6rd Tunnel MTU is configured with a high value so if the packet size is larger than the configured value, fragmentation occurs at the egress interface (towards the IPv4 access network).

---

## Examples: 6rd and 6to4 Configurations

- [Example: 6rd with Interface-Style Service Set Configuration on page 284](#)
- [Example: 6rd with Next-Hop-Style Service Set Configuration on page 285](#)
- [Example: 6rd Anycast Configuration on page 287](#)
- [Example: Hairpinning Between 6rd Domains Configuration on page 288](#)
- [Example: 6to4 Configuration on page 290](#)

### Example: 6rd with Interface-Style Service Set Configuration

```
chassis {
  fpc 0 {
    pic 0 {
      inline-services {
        bandwidth 10g;
      }
    }
  }
}
services {
  service-set vrf-intf-service-set {
    software-rules swire01-r1;
    interface-service {
      service-interface si-0/0/0.0;
    }
  }
  software {
    software-concentrator {
      v6rd swire01-rd1 {
        software-address 30.30.30.1;
        ipv4-prefix 10.10.0.0/16;
        v6rd-prefix 3040::0/16;
        mtu-v4 9192;
      }
    }
  }
  rule swire01-r1 {
    match-direction input;
    term t1 {
```

### Example: 6rd with Next-Hop-Style Service Set Configuration

---

Copyright © 2016, Juniper Networks, Inc. 285

```
services {
  service-set vrf-nh-service-set {
    software-rules swire01-r1;
    next-hop-service {
      inside-service-interface si-0/0/0.1;
      outside-service-interface si-0/0/0.2;
    }
  }
  software {
    software-concentrator {
      v6rd swire01-rd1 {
        software-address 30.30.30.1;
        ipv4-prefix 10.10.0.0/16;
        v6rd-prefix 3040::0/16;
        mtu-v4 9192;
      }
    }
    rule swire01-r1 {
      match-direction input;
      term t1 {
        then {
          v6rd swire01-rd1;
        }
      }
    }
  }
}
interfaces {
  si-0/0/0 {
    unit 1 {
      family inet;
      family inet6;
      service-domain inside;
    }
    unit 2 {
      family inet;
      family inet6;
      service-domain outside;
    }
  }
  ge-0/2/7 {
    unit 0 {
      family inet {
        address 10.10.10.1/16;
      }
    }
  }
  ge-0/2/8 {
    unit 0 {
      family inet6 {
        address 3abc::1/16;
      }
    }
  }
}
routing-instances {
```



```

v6rd-vrf {
  instance-type vrf;
  interface si-0/0/0.1;
  interface ge-0/2/7.0;
  route-distinguisher 1.1.1.1;
  vrf-target target:100:100;
}
}

```

### Example: 6rd Anycast Configuration

```

chassis {
  fpc 0 {
    pic 0 {
      inline-services {
        bandwidth 10g;
      }
    }
    pic 2 {
      inline-services {
        bandwidth 1g;
      }
    }
  }
}
services {
  service-set anycast-nh-set1 {
    software-rules swire01-r1;
    next-hop-service {
      inside-service-interface si-0/0/0.1;
      outside-service-interface si-0/0/0.2;
    }
  }
  service-set anycast-nh-set2 {
    software-rules swire01-r1;
    next-hop-service {
      inside-service-interface si-0/2/0.1;
      outside-service-interface si-0/2/0.2;
    }
  }
}
software {
  software-concentrator {
    v6rd swire01-rd1 {
      software-address 30.30.30.1;
      ipv4-prefix 10.10.0.0/16;
      v6rd-prefix 3040::0/16;
      mtu-v4 9192;
    }
  }
  rule swire01-r1 {
    match-direction input;
    term t1 {
      then {
        v6rd swire01-rd1;
      }
    }
  }
}

```

```

    }
  }
}
interfaces {
  si-0/0/0 {
    unit 0 {
      family inet;
      family inet6;
    }
    unit 1 {
      family inet;
      family inet6;
      service-domain inside;
    }
    unit 2 {
      family inet;
      family inet6;
      service-domain outside;
    }
  }
}
si-0/2/0 {
  unit 0 {
    family inet;
    family inet6;
  }
  unit 1 {
    family inet;
    family inet6;
    service-domain inside;
  }
  unit 2 {
    family inet;
    family inet6;
    service-domain outside;
  }
}
ge-0/2/7 {
  unit 0 {
    family inet {
      address 10.10.10.1/16;
    }
  }
}
ge-0/2/8 {
  unit 0 {
    family inet6 {
      address 3abc::1/16;
    }
  }
}
}
}

```

### Example: Hairpinning Between 6rd Domains Configuration

This example uses an interface service-set and a next-hop service set as hairpinning domains.

```

chassis {
  fpc 0 {
    pic 0 {
      inline-services {
        bandwidth 10g;
      }
    }
  }
}
services {
  service-set hairpin-intf-service-set {
    software-rules swire01-r1;
    interface-service {
      service-interface si-0/0/0.0;
    }
  }
  service-set hairpin-nh-service-set {
    software-rules swire01-r2;
    next-hop-service {
      inside-service-interface si-0/0/0.1;
      outside-service-interface si-0/0/0.2;
    }
  }
}
software {
  software-concentrator {
    v6rd swire01-rd1 {
      software-address 30.30.30.1;
      ipv4-prefix 10.10.0.0/16;
      v6rd-prefix 3040::0/16;
      mtu-v4 9192;
    }
    v6rd swire01-rd2 {
      software-address 60.60.60.1;
      ipv4-prefix 40.40.40.0/24;
      v6rd-prefix 3050::0/16;
      mtu-v4 9192;
    }
  }
  rule swire01-r1 {
    match-direction input;
    term t1 {
      then {
        v6rd swire01-rd1;
      }
    }
  }
  rule swire01-r2 {
    match-direction input;
    term t1 {
      then {
        v6rd swire01-rd2;
      }
    }
  }
}
}

```

```
interfaces {
  si-0/0/0 {
    unit 0 {
      family inet;
      family inet6;
    }
    unit 1 {
      family inet;
      family inet6;
      service-domain inside;
    }
    unit 2 {
      family inet;
      family inet6;
      service-domain outside;
    }
  }
  ge-0/2/7 {
    unit 0 {
      family inet {
        service {
          input {
            service-set hairpin-intf-service-set;
          }
          output {
            service-set hairpin-intf-service-set;
          }
        }
      }
      address 10.10.10.1/16;
    }
  }
  ge-0/2/8 {
    unit 0 {
      family inet {
        address 40.40.40.1/24;
      }
    }
  }
}
```

### Example: 6to4 Configuration

```
chassis {
  fpc 0 {
    pic 0 {
      inline-services {
        bandwidth 10g;
      }
    }
  }
}
services {
  service-set 6to4-intf-service-set {
    software-rules shenick01-r1;
    interface-service {
```

```

        service-interface si-0/0/0.0;
    }
}
interfaces {
    si-0/0/0 {
        unit 0 {
            family inet;
            family inet6;
        }
        unit 1 {
            family inet;
            family inet6;
            service-domain inside;
        }
        unit 2 {
            family inet;
            family inet6;
            service-domain outside;
        }
    }
}
ge-0/2/7 {
    unit 0 {
        family inet {
            service {
                input {
                    service-set 6to4-intf-service-set;
                }
                output {
                    service-set 6to4-intf-service-set;
                }
            }
            address 10.10.10.1/16;
        }
    }
}
ge-0/2/8 {
    unit 0 {
        family inet6 {
            address 3abc::1/16;
        }
    }
}
}

```

**Related Documentation**

- [Configuring a 6to4 Provider-Managed Tunnel](#)



## CHAPTER 22

# Monitoring and Troubleshooting Softwires

- [Ping and Traceroute for DS-Lite on page 293](#)
- [Monitoring Softwire Statistics on page 293](#)
- [Monitoring CGN, Stateful Firewall, and Softwire Flows on page 295](#)

## 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 4

# Enabling Traffic to Pass Securely Using ALGs

- [ALG Overview on page 299](#)
- [ALGs Configuration Overview on page 325](#)



## CHAPTER 23

# ALG Overview

- [ALG Descriptions on page 299](#)
- [ALGs Available by Default for Junos OS Address Aware NAT on page 321](#)

## 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 299](#)
- [ALG Support Details on page 300](#)
- [Juniper Networks Defaults on page 310](#)
- [Examples: Referencing the Preset Statement from the Junos OS Default Group on page 320](#)

## Supported ALGs

[Table 10 on page 299](#) lists ALGs supported by Junos OS. For information about which ALGs are supported on MS-DPCs, MS-MPCs, or MS-MICs, see [“ALGs Available by Default for Junos OS Address Aware NAT” on page 127](#).

**Table 10: 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

Table 10: ALGs Supported by Junos OS (*continued*)

ALGs Supported	v4 - v4	v4 - v6	v6 - v6	DS-Lite
ICMP	Yes	Yes	Yes	Yes
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 301](#)
- [Basic UDP ALG on page 302](#)
- [BOOTP on page 302](#)
- [DCE RPC Services on page 302](#)
- [DNS on page 302](#)
- [FTP on page 303](#)
- [H323 on page 303](#)
- [ICMP on page 304](#)

- IIOp on page 304
- IP on page 304
- NetBIOS on page 304
- NetShow on page 305
- ONC RPC Services on page 305
- PPTP on page 305
- RealAudio on page 305
- Sun RPC and RPC Portmap Services on page 306
- RTSP on page 307
- SIP on page 308
- SNMP on page 308
- SQLNet on page 308
- TFTP on page 309
- Traceroute on page 309
- UNIX Remote-Shell Services on page 309
- Winframe on page 310

---

### 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 when a flow 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 when a flow 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 System (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 closes the session only 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.

## IIOIP

---

The Oracle Application Server Name Server Internet Inter-ORB Protocol (IIOIP). This ALG is used in Common Object Request Broker Architecture (CORBA) based on distributed computing. Even though CORBA and IIOIP are Object Management Group (OMG) standards, no fixed port is assigned for IIOIP. 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 IIOIP be configured for TCP port 1975 for Java VM IIOIP, and 3075 for CORBA applications ORBIX, a CORBA framework from Iona Technologies.

## IP

---

The IP ALG is used to create unidirectional 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 NAPT, which causes 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 11 on page 305](#).

**Table 11: 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 11: 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 12 on page 306](#).

Table 12: Supported RPC Services

Name	Description	Comments
<b>rpc-mountd</b>	Network File Server (NFS) mount daemon; for details, see the UNIX man page for <b>rpc.mountd(8)</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050).
<b>rpc-nfsprog</b>	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).
<b>rpc-nisplus</b>	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).
<b>rpc-nlockmgr</b>	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, <b>rpc-nlockmgr</b> service can be allowed or blocked based on RPC program 100021.
<b>rpc-pcnfsd</b>	Kernel statistics server. For details, see the UNIX man pages for <b>rstatd</b> and <b>rpc.rstatd</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc-rstat</b> service can be allowed or blocked based on RPC program 150001.

Table 12: Supported RPC Services (*continued*)

Name	Description	Comments
<b>rpc-rwall</b>	Used to write a message to users; for details, see the UNIX man page for <b>rpc.rwalld</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc-rwall</b> service can be allowed or blocked based on RPC program 150008.
<b>rpc-ybind</b>	NIS binding process. For details, see the UNIX man page for <b>ybind</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc-ybind</b> service can be allowed or blocked based on RPC program 100007.
<b>rpc-yppasswd</b>	NIS password server. For details, see the UNIX man page for <b>yppasswd</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc-yppasswd</b> service can be allowed or blocked based on RPC program 100009.
<b>rpc-ybserv</b>	NIS server. For details, see the UNIX man page for <b>ybserv</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc-ybserv</b> service can be allowed or blocked based on RPC program 100004.
<b>rpc-ypupdated</b>	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, <b>rpc-ypupdated</b> service can be allowed or blocked based on RPC program 100028.
<b>rpc-ypxfrd</b>	NIS map transfer server. For details, see the UNIX man page for <b>rpc.ypxfrd</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc-ypxfrd</b> 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. SIP sessions on the MS-DPC have no time limit.

---

## 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 OS 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 OS defaults group.

To view the full set of available preset statements from the Junos OS default group, issue the **show groups junos-defaults** configuration mode command. The following example displays the list of Junos OS 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-cvspserver {
        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 325](#); for details about a specific protocol, see [“ALG Descriptions” on page 299](#).

### Examples: Referencing the Preset Statement from the Junos OS Default Group

The following example is a preset statement from the Junos OS 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 OS default statement from the Junos OS default groups, include the **junos-default-name** statement at the applicable hierarchy level. For example, to reference the Junos OS 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 when any other more specific application matches the same traffic, the IP ALG is not 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**
- [Configuring Application Sets on page 325](#)
  - [Configuring Application Protocol Properties on page 325](#)

## ALGs Available by Default for Junos OS Address Aware NAT

The following Application Level Gateways (ALGs) listed in [Table 9 on page 127](#) 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 13: ALGs Available by Default**

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
Basic TCP ALG	yes	yes	<b>NOTE:</b> Specific Junos OS 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	<b>NOTE:</b> TCP tracker performs limited integrity and validation checks for UDP.
BOOTP	yes	no	<ul style="list-style-type: none"> <li><b>junos-bootpc</b></li> <li><b>junos-bootps</b></li> </ul>
DCE RPC Services	yes	yes	<ul style="list-style-type: none"> <li><b>junos-dce-rpc-portmap</b></li> <li><b>junos-dcerpc-endpoint-mapper-service</b></li> <li><b>junos-dcerpc-msexchange-directory-nsp</b></li> <li><b>junos-dcerpc-msexchange-directory-rfr</b></li> <li><b>junos-dcerpc-msexchange-information-store</b></li> </ul>
DNS	yes	yes	<ul style="list-style-type: none"> <li><b>junos-dns-tcp</b></li> <li><b>junos-dns-udp</b></li> </ul>
FTP	yes	yes	<ul style="list-style-type: none"> <li><b>junos-ftp</b></li> </ul>
H323	yes	no	<ul style="list-style-type: none"> <li><b>junos-h323</b></li> </ul>
ICMP	yes	yes	<ul style="list-style-type: none"> <li><b>junos-icmp-all</b></li> <li><b>junos-icmp-ping</b></li> </ul>
IIOp	yes	no	<ul style="list-style-type: none"> <li><b>junos-iiop-java</b></li> <li><b>junos-iiop-orbix</b></li> </ul>

Table 13: ALGs Available by Default (*continued*)

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
IP	yes	The TCP tracker, available by default on these platforms, performs limited integrity and validation checks.	<ul style="list-style-type: none"> <li>• <b>junos-ip</b></li> </ul>
NETBIOS	yes	no	<ul style="list-style-type: none"> <li>• <b>junos-netbios-datagram</b></li> <li>• <b>junos-netbios-name-tcp</b></li> <li>• <b>junos-netbios-name-udp</b></li> <li>• <b>junos-netbios-session</b></li> </ul>
NETSHOW	yes	no	<ul style="list-style-type: none"> <li>• <b>junos-netshow</b></li> </ul>
PPTP	yes	yes	<ul style="list-style-type: none"> <li>• <b>junos-pptp</b></li> </ul>
REALAUDIO	yes	no	<ul style="list-style-type: none"> <li>• <b>junos-realaudio</b></li> </ul>
Sun RPC and RPC Port Map Services	yes	yes	<ul style="list-style-type: none"> <li>• <b>junos-rpc-portmap-tcp</b></li> <li>• <b>junos-rpc-portmap-udp</b></li> </ul>
RTSP	yes	yes	<ul style="list-style-type: none"> <li>• <b>junos-rtsp</b></li> </ul>
SIP	yes	Yes	<ul style="list-style-type: none"> <li>• <b>junos-sip</b></li> </ul> <p>The SIP <b>callid</b> is <i>not</i> translated in <b>register</b> messages.</p> <p><b>NOTE:</b> SIP sessions are limited to 12 hours (720 minutes) for NAT processing on the MS-MIC and MS-MPC interface cards. SIP sessions on the MS-DPC have no time limits.</p>
SNMP	yes	No	<ul style="list-style-type: none"> <li>• <b>junos-snmp-get</b></li> <li>• <b>junos-snmp-get-next</b></li> <li>• <b>junos-snmp-response junos-snmp-trap</b></li> </ul>
SQLNET	yes	yes	<ul style="list-style-type: none"> <li>• <b>junos-sqlnet</b></li> </ul>
TFTP	yes	yes	<ul style="list-style-type: none"> <li>• <b>junos-tftp</b></li> </ul>
Traceroute	yes	yes	<ul style="list-style-type: none"> <li>• <b>junos-traceroute</b></li> </ul>
Unix Remote Shell Service	yes	Yes	<ul style="list-style-type: none"> <li>• <b>junos-rsh</b></li> </ul>
WINFrame	yes	No	<ul style="list-style-type: none"> <li>• <b>junos-citrix-winframe</b></li> <li>• <b>junos-citrix-winframe-udp</b></li> </ul>
TALK-UDP	No	Yes	<ul style="list-style-type: none"> <li>• <b>junos-talk-udp</b></li> </ul>

Table 13: ALGs Available by Default *(continued)*

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
MS RPC	No	Yes	<ul style="list-style-type: none"><li>• junos-rpc-portmap-tcp</li><li>• junos-rpc-portmap-udp</li><li>• junos-rpc-services-tcp</li><li>• junos-rpc-services-udp</li></ul>

**Related Documentation**    • [ALG Descriptions on page 299](#)

# ALGs Configuration Overview

- [Configuring Application Sets on page 325](#)
- [Configuring Application Protocol Properties on page 325](#)
- [Examples: Configuring Application Protocols on page 343](#)
- [Verifying the Output of ALG Sessions on page 344](#)
- [ICMP, Ping, and Traceroute ALGs for MS-MICs and MS-MPCs on page 350](#)
- [Monitoring Port Control Protocol Operations on page 351](#)

## 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 343.

### Related Documentation

- [ALG Descriptions on page 299](#)
- [Configuring Application Protocol Properties on page 325](#)
- [Examples: Configuring Application Protocols on page 343](#)
- [Verifying the Output of ALG Sessions on page 344](#)

## 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 325.

This section includes the following tasks for configuring applications:

- [Configuring an Application Protocol on page 326](#)
- [Configuring the Network Protocol on page 328](#)
- [Configuring the ICMP Code and Type on page 329](#)
- [Configuring Source and Destination Ports on page 331](#)
- [Configuring the Inactivity Timeout Period on page 334](#)
- [Configuring SIP on page 334](#)
- [Configuring an SNMP Command for Packet Matching on page 342](#)
- [Configuring an RPC Program Number on page 342](#)
- [Configuring the TTL Threshold on page 342](#)
- [Configuring a Universal Unique Identifier on page 342](#)

## 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 14 on page 326](#) shows the list of supported protocols. For more information about specific protocols, see “[ALG Descriptions](#)” on page 299.

**Table 14: Application Protocols Supported by Services Interfaces**

Protocol Name	CLI Value	Comments
Bootstrap protocol (BOOTP)	<b>bootp</b>	Supports BOOTP and dynamic host configuration protocol (DHCP).
Distributed Computing Environment (DCE) remote procedure call (RPC)	<b>dce-rpc</b>	Requires the <b>protocol</b> statement to have the value <b>udp</b> or <b>tcp</b> . Requires a <b>uuid</b> value. You cannot specify <b>destination-port</b> or <b>source-port</b> values.
DCE RPC portmap	<b>dce-rpc-portmap</b>	Requires the <b>protocol</b> statement to have the value <b>udp</b> or <b>tcp</b> . Requires a <b>destination-port</b> value.



**Table 14: Application Protocols Supported by Services Interfaces** (*continued*)

Protocol Name	CLI Value	Comments
Domain Name System (DNS)	<b>dns</b>	Requires the <b>protocol</b> statement to have the value <b>udp</b> . This application protocol closes the DNS flow as soon as the DNS response is received.
Exec	<b>exec</b>	Requires the <b>protocol</b> statement to have the value <b>tcp</b> or to be unspecified. Requires a <b>destination-port</b> value.
FTP	<b>ftp</b>	Requires the <b>protocol</b> statement to have the value <b>tcp</b> or to be unspecified. Requires a <b>destination-port</b> value.
H.323	<b>h323</b>	—
Internet Control Message Protocol (ICMP)	<b>icmp</b>	Requires the <b>protocol</b> statement to have the value <b>icmp</b> or to be unspecified.
Internet Inter-ORB Protocol	<b>iiop</b>	—
IP	<b>ip</b>	—
Login	<b>login</b>	—
NetBIOS	<b>netbios</b>	Requires the <b>protocol</b> statement to have the value <b>udp</b> or to be unspecified. Requires a <b>destination-port</b> value.
NetShow	<b>netshow</b>	Requires the <b>protocol</b> statement to have the value <b>tcp</b> or to be unspecified. Requires a <b>destination-port</b> value.
Point-to-Point Tunneling Protocol	<b>pptp</b>	—
RealAudio	<b>realaudio</b>	—
Real-Time Streaming Protocol (RTSP)	<b>rtsp</b>	Requires the <b>protocol</b> statement to have the value <b>tcp</b> or to be unspecified. Requires a <b>destination-port</b> value.
RPC User Datagram Protocol (UDP) or TCP	<b>rpc</b>	Requires the <b>protocol</b> statement to have the value <b>udp</b> or <b>tcp</b> . Requires a <b>rpc-program-number</b> value. You cannot specify <b>destination-port</b> or <b>source-port</b> values.
RPC port mapping	<b>rpc-portmap</b>	Requires the <b>protocol</b> statement to have the value <b>udp</b> or <b>tcp</b> . Requires a <b>destination-port</b> value.
Shell	<b>shell</b>	Requires the <b>protocol</b> statement to have the value <b>tcp</b> or to be unspecified. Requires a <b>destination-port</b> value.
Session Initiation Protocol	<b>sip</b>	—
SNMP	<b>snmp</b>	Requires the <b>protocol</b> statement to have the value <b>udp</b> or to be unspecified. Requires a <b>destination-port</b> value.

Table 14: Application Protocols Supported by Services Interfaces (*continued*)

Protocol Name	CLI Value	Comments
SQLNet	<b>sqlnet</b>	Requires the <b>protocol</b> statement to have the value <b>tcp</b> or to be unspecified. Requires a <b>destination-port</b> or <b>source-port</b> value.
Talk Program	<b>talk</b>	
Trace route	<b>traceroute</b>	Requires the <b>protocol</b> statement to have the value <b>udp</b> or to be unspecified. Requires a <b>destination-port</b> value.
Trivial FTP (TFTP)	<b>tftp</b>	Requires the <b>protocol</b> statement to have the value <b>udp</b> or to be unspecified. Requires a <b>destination-port</b> value.
WinFrame	<b>winframe</b>	—



**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 127](#)

## 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 15 on page 328](#) shows the list of the supported protocols.

Table 15: Network Protocols Supported by Services Interfaces

Network Protocol Type	CLI Value	Comments
IP Security (IPsec) authentication header (AH)	<b>ah</b>	—

Table 15: Network Protocols Supported by Services Interfaces (*continued*)

Network Protocol Type	CLI Value	Comments
External Gateway Protocol (EGP)	<b>egp</b>	—
IPsec Encapsulating Security Payload (ESP)	<b>esp</b>	—
Generic routing encapsulation (GR)	<b>gre</b>	—
ICMP	<b>icmp</b>	Requires an <b>application-protocol</b> value of <b>icmp</b> .
ICMPv6	<b>icmp6</b>	Requires an <b>application-protocol</b> value of <b>icmp</b> .
Internet Group Management Protocol (IGMP)	<b>igmp</b>	—
IP in IP	<b>ipip</b>	—
OSPF	<b>ospf</b>	—
Protocol Independent Multicast (PIM)	<b>pim</b>	—
Resource Reservation Protocol (RSVP)	<b>rsvp</b>	—
TCP	<b>tcp</b>	Requires a <b>destination-port</b> or <b>source-port</b> value unless you specify <b>application-protocol rcp</b> or <b>dce-rcp</b> .
UDP	<b>udp</b>	Requires a <b>destination-port</b> or <b>source-port</b> value unless you specify <b>application-protocol rcp</b> or <b>dce-rcp</b> .

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 16 on page 330](#) shows the list of supported ICMP values.

**Table 16: ICMP Codes and Types Supported by Services Interfaces**

CLI Statement	Description
<b>icmp-code</b>	<p>This value or keyword provides more specific information than <b>icmp-type</b>. Because the value's meaning depends upon the associated <b>icmp-type</b> value, you must specify <b>icmp-type</b> along with <b>icmp-code</b>. 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: <b>ip-header-bad</b> (0), <b>required-option-missing</b> (1)</p> <p>redirect: <b>redirect-for-host</b> (1), <b>redirect-for-network</b> (0), <b>redirect-for-tos-and-host</b> (3), <b>redirect-for-tos-and-net</b> (2)</p> <p>time-exceeded: <b>ttl-eq-zero-during-reassembly</b> (1), <b>ttl-eq-zero-during-transit</b> (0)</p> <p>unreachable: <b>communication-prohibited-by-filtering</b> (13), <b>destination-host-prohibited</b> (10), <b>destination-host-unknown</b> (7), <b>destination-network-prohibited</b> (9), <b>destination-network-unknown</b> (6), <b>fragmentation-needed</b> (4), <b>host-precedence-violation</b> (14), <b>host-unreachable</b> (1), <b>host-unreachable-for-TOS</b> (12), <b>network-unreachable</b> (0), <b>network-unreachable-for-TOS</b> (11), <b>port-unreachable</b> (3), <b>precedence-cutoff-in-effect</b> (15), <b>protocol-unreachable</b> (2), <b>source-host-isolated</b> (8), <b>source-route-failed</b> (5)</p>
<b>icmp-type</b>	<p>Normally, you specify this match in conjunction with the <b>protocol</b> 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): <b>echo-reply</b> (0), <b>echo-request</b> (8), <b>info-reply</b> (16), <b>info-request</b> (15), <b>mask-request</b> (17), <b>mask-reply</b> (18), <b>parameter-problem</b> (12), <b>redirect</b> (5), <b>router-advertisement</b> (9), <b>router-solicit</b> (10), <b>source-quench</b> (4), <b>time-exceeded</b> (11), <b>timestamp</b> (13), <b>timestamp-reply</b> (14), or <b>unreachable</b> (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 14 on page 326](#).

You can specify either a numeric value or one of the text synonyms listed in [Table 17 on page 331](#).

**Table 17: 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 17: 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 17: 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 17: 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*.

## 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 341](#)

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 335](#).



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 326](#). 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 336](#)
- [Incoming Calls on page 337](#)
- [Forwarded Calls on page 337](#)
- [Call Termination on page 337](#)
- [Call Re-INVITE Messages on page 337](#)
- [Call Session Timers on page 338](#)
- [Call Cancellation on page 338](#)
- [Forking on page 338](#)
- [SIP Messages on page 338](#)
- [SIP Headers on page 338](#)
- [SIP Body on page 341](#)

### ***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 18 on page 339 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 18: 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 18: 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.

**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 326](#).

## 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 326](#).

## 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 326](#).

## 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 326. 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 299](#)
- [Configuring Application Sets on page 325](#)
- [Configuring Application Protocol Properties on page 325](#)
- [Verifying the Output of ALG Sessions on page 344](#)

## Verifying the Output of ALG Sessions

This section contains examples of successful output from ALG sessions and information on system log configuration. You can compare the results of your sessions to check whether the configurations are functioning correctly.

- [FTP Example on page 344](#)
- [RTSP ALG Example on page 346](#)
- [System Log Messages on page 349](#)

### FTP Example

This example analyzes the output during an active FTP session. It consists of four different flows; two are control flows and two are data flows. The example consists of the following parts:

- [Sample Output on page 344](#)
- [FTP System Log Messages on page 345](#)
- [Analysis on page 345](#)
- [Troubleshooting Questions on page 346](#)

### Sample Output

The following is a complete sample output from the **show services stateful-firewall conversations application-protocol ftp** operational mode command:

```
user@host>show services stateful-firewall conversations application-protocol ftp
Interface: ms-1/3/0, Service set: CLBJI1-AAF001
Conversation: ALG protocol: ftp
  Number of initiators: 2, Number of responders: 2
Flow
TCP      1.1.79.2:14083 ->      2.2.2.2:21      Watch    I      Frm count
      NAT source      1.1.79.2:14083 ->      194.250.1.237:50118
TCP      1.1.79.2:14104 ->      2.2.2.2:20      Forward  I      3
      NAT source      1.1.79.2:14104 ->      194.250.1.237:50119
TCP      2.2.2.2:21 ->      194.250.1.237:50118 Watch    O      12
      NAT dest      194.250.1.237:50118 ->      1.1.79.2:14083
TCP      2.2.2.2:20 ->      194.250.1.237:50119 Forward  O      5
      NAT dest      194.250.1.237:50119 ->      1.1.79.2:14104
```

For each flow, the first line shows flow information, including protocol (TCP), source address, source port, destination address, destination port, flow state, direction, and frame count.

- The state of a flow can be **Watch**, **Forward**, or **Drop**:
  - A **Watch** flow state indicates that the control flow is monitored by the ALG for information in the payload. NAT processing is performed on the header and payload as needed.
  - A **Forward** flow forwards the packets without monitoring the payload. NAT is performed on the header as needed.

- A **Drop** flow drops any packet that matches the 5 tuple.
- The frame count (**Frm count**) shows the number of packets that were processed on that flow.

The second line shows the NAT information.

- **source** indicates source NAT.
- **dest** indicates destination NAT.
- The first address and port in the NAT line are the original address and port being translated for that flow.
- The second address and port in the NAT line are the translated address and port for that flow.

### FTP System Log Messages

System log messages are generated during an FTP session. For more information about system logs, see [“System Log Messages” on page 349](#).

The following system log messages are generated during creation of the FTP control flow:

- Rule Accept system log:  
Oct 27 11:42:54 (FPC Slot 1, PIC Slot 1) {ss\_ftp}[FWNAT]: ASP\_SFW\_RULE\_ACCEPT: proto 6 (TCP) application: ftp, fe-3/3/3.0:1.1.1.2:4450 -> 2.2.2.2:21, Match SFW accept rule-set:, rule: ftp, term: 1
- Create Accept Flow system log:  
Oct 27 11:42:54 (FPC Slot 1, PIC Slot 1) {ss\_ftp}[FWNAT]: ASP\_SFW\_CREATE\_ACCEPT\_FLOW: proto 6 (TCP) application: ftp, fe-3/3/3.0:1.1.1.2:4450 -> 2.2.2.2:21, creating forward or watch flow
- System log for data flow creation:  
Oct 27 11:43:30 (FPC Slot 1, PIC Slot 1) {ss\_ftp}[FWNAT]: ASP\_SFW\_FTP\_ACTIVE\_ACCEPT: proto 6 (TCP) application: ftp, so-2/1/2.0:2.2.2.2:20 -> 1.1.1.2:50726, Creating FTP active mode forward flow

### Analysis

#### Control Flows

The control flows are established after the three-way handshake is complete.

- Control flow from FTP client to FTP server. TCP destination port is 21.

```
TCP          1.1.79.2:14083 ->      2.2.2.2:21    Watch    I
13
NAT source   1.1.79.2:14083  ->    194.250.1.237:50118
```

- Control flow from FTP server to FTP client. TCP source port is 21.

```

TCP          2.2.2.2:21    -> 194.250.1.237:50118 Watch    0
12
NAT dest     194.250.1.237:50118 ->      1.1.79.2:14083

```

### Data Flows

A data port of 20 is negotiated for data transfer during the course of the FTP control protocol. These two flows are data flows between the FTP client and the FTP server:

```

TCP          1.1.79.2:14104 ->      2.2.2.2:20    Forward I          3
NAT source   1.1.79.2:14104 -> 194.250.1.237:50119
TCP          2.2.2.2:20    -> 194.250.1.237:50119 Forward O          5
NAT dest     194.250.1.237:50119 ->      1.1.79.2:14104

```

### Troubleshooting Questions

- How do I know if the FTP ALG is active?
  - The ALG protocol field in the conversation should display **ftp**.
  - There should be a valid frame count (**Frm count**) in the control flows.
  - A valid frame count in the data flows indicates that data transfer has taken place.
- What do I need to check if the FTP connection is established but data transfer does not take place?
  - Most probably, the control connection is up, but the data connection is down.
  - Check the conversations output to determine whether both the control and data flows are present.
- How do I interpret each flow? What does each flow mean?
  - FTP control flow initiator flow—Flow with destination port 21
  - FTP control flow responder flow—Flow with source port ;21
  - FTP data flow initiator flow—Flow with destination port 20
  - FTP data flow responder flow—Flow with source port 20

### RTSP ALG Example

The following is an example of an RTSP conversation. The application uses the RTSP protocol for control connection. Once the connection is set up, the media is sent using UDP protocol (RTP).

This example consists of the following:

- [Sample Output on page 347](#)
- [Analysis on page 347](#)
- [Troubleshooting Questions on page 347](#)

## Sample Output

Here is the output from the **show services stateful-firewall conversations** operational mode command:

```
user@host# show services stateful-firewall conversations
Interface: ms-3/2/0, Service set: svc_set
Conversation: ALG protocol: rtsp
Number of initiators: 5, Number of responders: 5
```

Flow			State	Dir	Frm count
TCP	1.1.1.3:58795	-> 2.2.2.2:554	Watch	I	7
UDP	1.1.1.3:1028	-> 2.2.2.2:1028	Forward	I	0
UDP	1.1.1.3:1029	-> 2.2.2.2:1029	Forward	I	0
UDP	1.1.1.3:1030	-> 2.2.2.2:1030	Forward	I	0
UDP	1.1.1.3:1031	-> 2.2.2.2:1031	Forward	I	0
TCP	2.2.2.2:554	-> 1.1.1.3:58795	Watch	O	5
UDP	2.2.2.2:1028	-> 1.1.1.3:1028	Forward	O	6
UDP	2.2.2.2:1029	-> 1.1.1.3:1029	Forward	O	0
UDP	2.2.2.2:1030	-> 1.1.1.3:1030	Forward	O	3
UDP	2.2.2.2:1031	-> 1.1.1.3:1031	Forward	O	0

## Analysis

An RTSP conversation should consist of TCP flows corresponding to the RTSP control connection. There should be two flows, one in each direction, from client to server and from server to client:

TCP	1.1.1.3:58795	-> 2.2.2.2:554	Watch	I	7
TCP	2.2.2.2:554	-> 1.1.1.3:58795	Watch	O	5

- The RTSP control connection for the initiator flow is sent from destination port 554.
- The RTSP control connection for the responder flow is sent from source port 554.

The UDP flows correspond to RTP media sent over the RTSP connection.

## Troubleshooting Questions

- Media does not work when the RTSP ALG is configured. What do I do?
  - Check RTSP conversations to see whether both TCP and UDP flows exist.
  - The ALG protocol should be displayed as **rtsp**.



**NOTE:** The state of the flow is displayed as **Watch**, because the ALG processing is taking place and the client is essentially “watching” or processing payload corresponding to the application. For FTP and RTSP ALG flows, the control connections are always **Watch** flows.

- How do I check for ALG errors?
  - You can check for errors by issuing the following command. Each ALG has a separate field for ALG packet errors.

```
user@host# show services stateful-firewall statistics extensive
```

```
Interface: ms-3/2/0
Service set: svc_set
New flows:
  Accepts: 1347, Discards: 0, Rejects: 0
Existing flows:
  Accepts: 144187, Discards: 0, Rejects: 0
Drops:
  IP option: 0, TCP SYN defense: 0
  NAT ports exhausted: 0
Errors:
  IP: 0, TCP: 276
  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
  Unknown: 0
TCP errors:
  TCP header length inconsistencies: 0
  Source or destination port number is zero: 0
  Illegal sequence number and flags combinations: 0
  SYN attack (multiple SYN messages seen for the same flow): 276
  First packet not a SYN message: 0
  TCP port scan (TCP handshake, RST seen from server for SYN): 0
  Bad SYN cookie response: 0
UDP errors:
  IP data length less than minimum UDP header length (8 bytes): 0
  Source or destination port number 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
ALG errors:
  BOOTP: 0, DCE-RPC: 0, DCE-RPC portmap: 0
  DNS: 0, Exec: 0, FTP: 0
  ICMP: 0
  Login: 0, NetBIOS: 0, NetShow: 0
  RPC: 0, RPC portmap: 0
  RTSP: 0, Shell: 0
  SNMP: 0, SQLNet: 0, TFTP: 0
  Traceroute: 0
```

## System Log Messages

Enabling system log generation and checking the system log are also helpful for ALG flow analysis. This section contains the following:

- [System Log Configuration on page 349](#)
- [System Log Output on page 349](#)

### System Log Configuration

You can configure the enabling of system log messages at a number of different levels in the Junos OS CLI. As shown in the following sample configurations, the choice of level depends on how specific you want the event logging to be and what options you want to include. For details on the configuration options, see the *Junos OS Administration Library for Routing Devices* (system level) or the *Junos OS Services Interfaces Library for Routing Devices* (all other levels).

1. At the topmost global level:

```
user@host# show system syslog
file messages {
    any any;
}
```

2. At the service set level:

```
user@host# show services service-set svc_set
syslog {
    host local {
        services any;
    }
}
stateful-firewall-rules allow_rtsp;
interface-service {
    service-interface ms-3/2/0;
}
```

3. At the service rule level:

```
user@host# show services stateful-firewall rule allow_rtsp
match-direction input-output;
term 0 {
    from {
        applications junos-rtsp;
    }
    then {
        accept;
        syslog;
    }
}
```

### System Log Output

System log messages are generated during flow creation, as shown in the following examples:

The following system log message indicates that the ASP matched an accept rule:

```
Oct 25 16:11:37 (FPC Slot 3, PIC Slot 2) {svc_set}[FWNAT]: ASP_SFW_RULE_ACCEPT:
proto 6 (TCP) application: rtsp, ge-2/0/1.0:1.1.1.2:35595 -> 2.2.2.2:554, Match SFW accept
rule-set: , rule: allow_rtsp, term: 0
```

For a complete listing of system log messages, see the [System Log Explorer](#).

**Related  
Documentation**

- [ALG Descriptions on page 299](#)
- [Configuring Application Sets on page 325](#)
- [Configuring Application Protocol Properties on page 325](#)
- [Examples: Configuring Application Protocols on page 343](#)

---

## 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 299](#)

## Monitoring Port Control Protocol Operations

You can monitor Port Control Protocol (PCP) operations with the following operational commands:

- **show services nat mappings pcsp**
- **show services nat mappings endpoint-independent**
- **show services pcsp statistics protocol**

The following are examples of the output of these commands.

```
user@host> show services nat mappings pcsp
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 pcsp 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 5

# Securing Content Using Junos Network Secure and IDS

- [Junos Network Secure Overview on page 355](#)
- [Junos Network Secure Configuration Overview on page 359](#)
- [IDS Configuration Overview on page 383](#)
- [Monitoring Junos Network Secure on page 397](#)



# Junos Network Secure Overview

- [Junos Network Secure Overview on page 355](#)

## Junos Network Secure Overview

---

Routers use firewalls to track and control the flow of traffic. Adaptive Services and MultiServices PICs employ a type of firewall called a *stateful firewall*. 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 359](#).

## Stateful Firewall Support for Application Protocols

By inspecting the application protocol data, the AS or MultiServices PIC 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 359](#)
- [Configuring Stateful Firewall Rule Sets on page 363](#)
- [Examples: Configuring Stateful Firewall Rules on page 363](#)
- [Example: BOOTP and Broadcast Addresses on page 367](#)
- [Example: Configuring Layer 3 Services and the Services SDK on Two PICs on page 367](#)
- [Example: Virtual Routing and Forwarding \(VRF\) and Service Configuration on page 379](#)

## 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 360](#)
- [Configuring Match Conditions in Stateful Firewall Rules on page 360](#)
- [Configuring Actions in Stateful Firewall Rules on page 362](#)

## 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 9](#).

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 363.

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 325.

- 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 19 on page 363](#) 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 19: 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.



**NOTE:** You can define the service-set and assign it either as interface style or next-hop style.

#### Related Documentation

- [Example: BOOTP and Broadcast Addresses on page 367](#)
- [Example: Dynamic Source NAT as a Next-Hop Service on page 124](#)
- [Example: Virtual Routing and Forwarding \(VRF\) and Service Configuration on page 379](#)
- [Example: Service Interfaces Configuration](#)
- [Example: Configuring Layer 3 Services and the Services SDK on Two PICs on page 367](#)



- [Configuring Service Sets to be Applied to Services Interfaces on page 9](#)

## 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;
                    }
                }
            }
        }
    }
}
```



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 ACL rules.

```
[edit services service-set App-Aware-Set]
user@host# set aacl-rules rule-name
```

In this example, the ACL 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;
}
acl-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 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-aacl
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-aacl;
                    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;
  }
}
}
```





# IDS Configuration Overview

- [Understanding SYN Cookie Protection on page 383](#)
- [Configuring IDS Rules on page 384](#)
- [Configuring IDS Rule Sets on page 392](#)
- [Examples: Configuring IDS Rules on page 392](#)

## Understanding SYN Cookie Protection

---

SYN cookie is a stateless SYN proxy mechanism you can use in conjunction with other defenses against a SYN flood attack.

As with traditional SYN proxying, SYN cookie is activated when the SYN flood attack threshold is exceeded. However, because SYN cookie is stateless, it does not set up a session or policy and route lookups upon receipt of a SYN segment, and it maintains no connection request queues. This dramatically reduces CPU and memory usage and is the primary advantage of using SYN cookie over the traditional SYN proxying mechanism.

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.

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.



**NOTE:** The use of SYN cookie or SYN proxy enables the router device to protect the TCP servers behind it from SYN flood attacks in IPv6 flows.

---

Related Documentation • [Configuring IDS Rule Sets on page 392](#)

## 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 359](#).

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 385](#)
- [Configuring Match Conditions in IDS Rules on page 386](#)
- [Configuring Actions in IDS Rules on page 387](#)

## 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 9](#).

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 363.

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 325.

- 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 not 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 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 392](#)
  - [Examples: Configuring IDS Rules on page 392](#)

---

## 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 384](#)
  - [Examples: Configuring IDS Rules on page 392](#)

---

## 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 384](#)
  - [Configuring IDS Rule Sets on page 392](#)



## CHAPTER 28

# Monitoring Junos Network Secure

- [Monitoring Stateful Firewall Conversations on page 397](#)
- [Monitoring CGN, Stateful Firewall, and Software Flows on page 397](#)
- [Monitoring Global Stateful Firewall Statistics on page 398](#)

### 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 6

# Creating Secure Tunnels Using Junos VPN Site Secure

- [Junos VPN Site Secure Overview on page 401](#)
- [Junos VPN Site Secure Configuration Overview on page 413](#)
- [Enhancing Security with Static IPsec over VRF on page 487](#)
- [Dynamically Assigning Tunnels Using Junos VPN Site Secure on page 495](#)
- [Enabling IPsec for the Services SDK on page 549](#)



## CHAPTER 29

# Junos VPN Site Secure Overview

- [Understanding Junos VPN Site Secure on page 401](#)
- [Authentication Algorithms on page 404](#)
- [Encryption Algorithms on page 404](#)
- [IPsec Protocols on page 406](#)
- [Supported IPsec and IKE Standards on page 408](#)
- [IPsec Terms and Acronyms on page 409](#)

## 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 401](#)
- [Security Associations on page 402](#)
- [IKE on page 402](#)
- [Comparison of IPsec on ES PICs and Junos VPN Site Secure on Multiservices Line Cards on page 402](#)

## 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 20 on page 403](#) 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 20: 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 404](#)
- [Encryption Algorithms on page 404](#)
- [IPsec Protocols on page 406](#)
- [Service Sets for IPsec Tunnels on page 460](#)
- [Configuring Security Associations on page 415](#)
- [IPsec Hierarchy Level on page 706](#)

## 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 401](#)
- [Encryption Algorithms on page 404](#)

## 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 401](#)
- [Configuring IKE Proposals on page 435](#)
- [Configuring IPsec Proposals on page 445](#)
- [encryption on page 767](#)

## 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 21 on page 406](#).



**NOTE:** AH is not supported on the T Series, M120, and M320 routers.

**Figure 21: 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 the original IP header, AH header, and TCP header			

IPv4 packet after AH tunnel mode is applied

New IP header	AH header	Original IP header	TCP header	Data
Authenticating the original IP header, AH header, and TCP header				

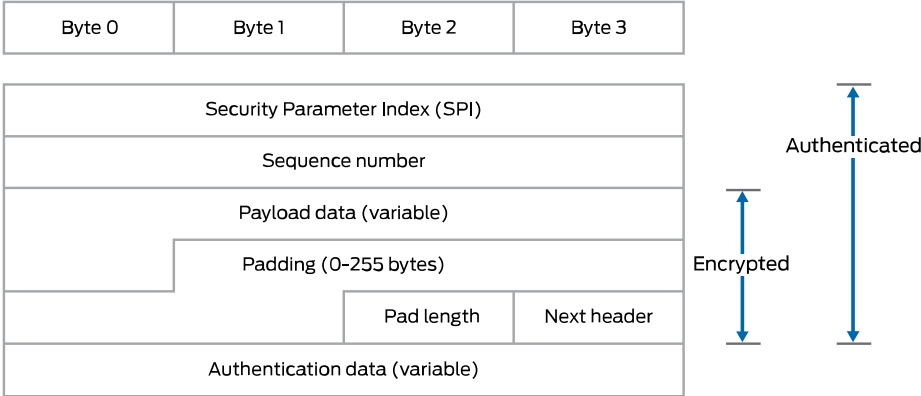
g015522



- 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 22 on page 407](#).

Figure 22: ESP Protocol


Header format




Original IPv4 packet before ESP is applied



IPv4 packet after ESP transport mode is applied



IPv4 packet after ESP tunnel mode is applied



g015521

- 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 401](#)
- [Configuring IPsec Proposals on page 445](#)
- [Configuring Security Associations on page 415](#)

- [protocol \(IPSec\) on page 859](#)

---

## 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

<b>Adaptive Services PIC</b>	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.
<b>Advanced Encryption Standard (AES)</b>	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 30

# Junos VPN Site Secure Configuration Overview

- [Minimum Security Association Configurations on page 413](#)
- [Configuring Security Associations on page 415](#)
- [Example: Configuring Manual SAs on page 421](#)
- [Configuring IKE Proposals on page 435](#)
- [Configuring IKE Policies on page 439](#)
- [Configuring IPsec Proposals on page 445](#)
- [Configuring IPsec Policies on page 450](#)
- [Configuring IPsec Rules on page 452](#)
- [Configuring IPsec Rule Sets on page 459](#)
- [Service Sets on page 460](#)
- [Configuring IPsec Service Sets on page 460](#)
- [Tracing Junos VPN Site Secure Operations on page 466](#)
- [Multitask Example: Configuring IPsec Services on page 468](#)
- [Example: Configuring Junos VPN Site Secure on MS MIC and MS-MPC on page 475](#)

## 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 413](#)
- [Minimum Dynamic SA Configuration on page 414](#)

### 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 401](#)
  - [Configuring Security Associations on page 415](#)
  - [Configuring IKE Proposals on page 435](#)
  - [Configuring IKE Policies on page 439](#)
  - [Configuring IPsec Proposals on page 445](#)
  - [Configuring IPsec Policies on page 450](#)

## 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 415](#)
- [Configuring Dynamic Security Associations on page 420](#)
- [Clearing Security Associations on page 420](#)

### 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 416](#)
- [Configuring the Protocol for a Manual IPsec SA on page 417](#)
- [Configuring the Security Parameter Index on page 417](#)
- [Configuring the Auxiliary Security Parameter Index on page 418](#)
- [Configuring Authentication for a Manual IPsec SA on page 418](#)
- [Configuring Encryption for a Manual IPsec SA on page 419](#)

### 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 450](#)
- [Configuring IPsec Proposals on page 445](#)

- [Configuring IKE Policies on page 439](#)
- [Configuring IKE Proposals on page 435](#)

## 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 421](#)
- [Overview and Topology on page 421](#)
- [Configuration on page 422](#)
- [Verification on page 433](#)

### 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

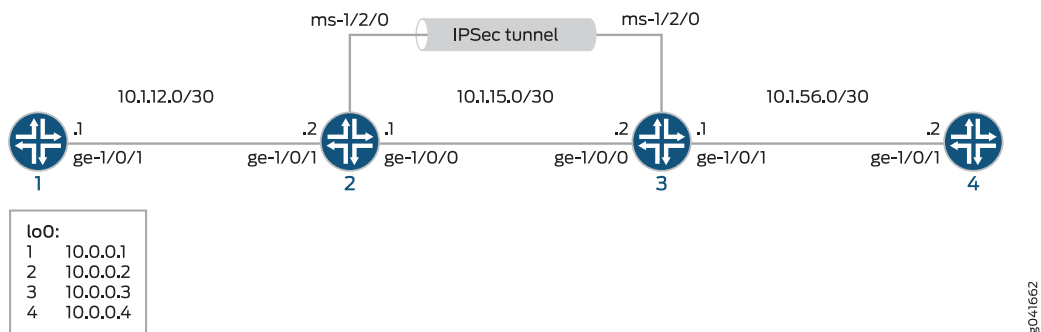
A security association (SA) is a simplex connection that enables two hosts to securely communicate with each other by means of IPsec. There are two types of SAs: manual SA and dynamic SA. This example explains a manual SA configuration.

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 23 on page 422](#) shows an IPsec topology that contains a group of four routers: Routers 1, 2, 3, and 4.

Figure 23: 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 422](#)
- [Configuring Router 2 on page 424](#)
- [Configuring Router 3 on page 428](#)
- [Configuring Router 4 on page 432](#)

### 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 10.1.12.1/30
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 10.1.12.1/30
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 10.1.12.1/30;
      }
    }
  }
  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 10.1.12.2/30
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 10.0.0.2/32
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 10.1.12.2/30
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 10.1.12.2/30;
      }
    }
  }
  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 10.0.0.2/32;
      }
    }
  }
}
```

```

    }
  }
}
...
}

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 "$ABC123"; ## SECRET-DATA
              }
              encryption {
                algorithm des-cbc;
                key ascii-text "$ABC123"; ## 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 "$ABC123"; ## SECRET-DATA
              }
              encryption {
                algorithm des-cbc;
                key ascii-text "$ABC123"; ## 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 433](#)
- [Verifying the Security Associations on Router 2 on page 434](#)
- [Verifying the Security Associations on Router 3 on page 434](#)

### 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 401](#)
  - [Configuring Security Associations on page 415](#)
  - [Example: Configuring IKE Dynamic SAs on page 509](#)

## 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 436](#)
- [Configuring the Authentication Method for an IKE Proposal on page 436](#)
- [Configuring the Diffie-Hellman Group for an IKE Proposal on page 437](#)
- [Configuring the Encryption Algorithm for an IKE Proposal on page 438](#)
- [Configuring the Lifetime for an IKE SA on page 438](#)
- [Example: Configuring an IKE Proposal on page 439](#)

## 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 445](#)
- [Configuring IKE Policies on page 439](#)
- [Configuring IPsec Policies on page 450](#)
- [Configuring Security Associations on page 415](#)

## 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 441](#)
- [Configuring the Mode for an IKE Policy on page 441](#)
- [Configuring the Proposals in an IKE Policy on page 441](#)
- [Configuring the Preshared Key for an IKE Policy on page 441](#)
- [Configuring the Local Certificate for an IKE Policy on page 442](#)
- [Configuring the Description for an IKE Policy on page 443](#)
- [Configuring Local and Remote IDs for IKE Phase 1 Negotiation on page 443](#)
- [Enabling Invalid SPI Recovery on page 444](#)
- [Example: Configuring an IKE Policy on page 444](#)

## 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 443](#)

## 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 495](#)
- [Configuring IKE Proposals on page 435](#)
- [Configuring IPsec Policies on page 450](#)
- [Configuring IPsec Proposals on page 445](#)
- [Configuring Security Associations on page 415](#)

## 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 446](#)
- [Configuring the Description for an IPsec Proposal on page 448](#)
- [Configuring the Encryption Algorithm for an IPsec Proposal on page 448](#)
- [Configuring the Lifetime for an IPsec SA on page 448](#)
- [Configuring the Protocol for a Dynamic SA on page 449](#)

## 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 450](#)
- [Configuring IKE Proposals on page 435](#)
- [Configuring IKE Policies on page 439](#)
- [Configuring Security Associations on page 415](#)

## 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 450](#)
- [Configuring Perfect Forward Secrecy on page 451](#)
- [Configuring the Proposals in an IPsec Policy on page 451](#)
- [IPsec Policy for Dynamic Endpoints on page 451](#)
- [Example: Configuring an IPsec Policy on page 452](#)

### 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 445](#)
  - [Configuring IKE Proposals on page 435](#)
  - [Configuring IKE Policies on page 439](#)
  - [Configuring Security Associations on page 415](#)

---

## 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 454](#)
- [Configuring Match Conditions in IPsec Rules on page 454](#)
- [Configuring Actions in IPsec Rules on page 456](#)

## 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 457](#)
- [Configuring Destination Addresses for Dead Peer Detection on page 457](#)
- [Configuring or Disabling IPsec Anti-Replay on page 458](#)
- [Enabling System Log Messages on page 459](#)
- [Specifying the MTU for IPsec Tunnels on page 459](#)

### 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 459](#)
- [Configuring Security Associations on page 415](#)

## 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 452](#)
- [Configuring Security Associations on page 415](#)

---

## Service Sets

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 401](#)
- [Configuring Junos VPN Site Secure or IPsec VPN on page 549](#)

---

## 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 461](#)
- [Configuring IKE Access Profiles for IPsec Service Sets on page 462](#)
- [Configuring Certification Authorities for IPsec Service Sets on page 463](#)
- [Configuring or Disabling Antireplay Service on page 463](#)
- [Clearing the Don't-Fragment Bit on page 464](#)
- [Configuring Passive-Mode Tunneling on page 465](#)
- [Configuring the Tunnel MTU Value on page 466](#)

## 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" on page 452](#).



**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](#)” on page 452.

## 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 [“Tracing Junos VPN Site Secure Operations” on page 466](#). 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 7](#)
- [Configuring Service Sets to be Applied to Services Interfaces on page 9](#)
- [Configuring Service Set Limitations on page 15](#)
- [Configuring System Logging for Service Sets on page 26](#)

---

## 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 467](#)
- [Tracing IPsec PKI Operations on page 468](#)

## 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 439](#)
- [Configuring IKE Proposals on page 435](#)

---

## 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 469](#)
2. [Configuring the IKE Policy \(and Referencing the IKE Proposal\) on page 469](#)
3. [Configuring the IPsec Proposal on page 470](#)
4. [Configuring the IPsec Policy \(and Referencing the IPsec Proposal\) on page 471](#)
5. [Configuring the IPsec Rule \(and Referencing the IKE and IPsec Policies\) on page 471](#)

6. [Configuring IPsec Trace Options on page 473](#)
7. [Configuring the Access Profile \(and Referencing the IKE and IPsec Policies\) on page 473](#)
8. [Configuring the Service Set \(and Referencing the IKE Profile and the IPsec Rule\) on page 474](#)

## 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” on page 435](#).

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” on page 439](#).

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” on page 445](#).

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” on page 450](#).

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” on page 452](#).

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 Junos VPN Site Secure Operations” on page 466](#).

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 460](#).

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 on page 435](#)
- [Configuring IKE Policies on page 439](#)
- [Configuring IPsec Proposals on page 445](#)
- [Configuring IPsec Policies on page 450](#)
- [Configuring IPsec Rules on page 452](#)
- [Tracing Junos VPN Site Secure Operations on page 466](#)
- [Configuring an IKE Access Profile](#)
- [Configuring IPsec Service Sets on page 460](#)

## Example: Configuring Junos VPN Site Secure on MS MIC and MS-MPC



**NOTE:** You can follow the same procedure and use the same configuration given in this example, to configure Junos VPN Site Secure (previously known as IPsec features) on MS-MPCs.

This example contains the following sections:

- [Requirements on page 475](#)
- [Overview on page 476](#)
- [Configuration on page 476](#)
- [Verification on page 484](#)

### Requirements

This example uses the following hardware and software components:

- Two MX Series routers with MS-MICs
- Junos OS Release 13.2 or later

## Overview

Junos OS, Release 13.2, extends support for Junos VPN Site Secure (formerly known as IPsec features) to the newly-introduced Multiservices MIC and MPC (MS-MIC and MS-MPC) on MX Series routers. The Junos OS extension-provider packages come preinstalled and preconfigured on the MS-MIC and MS-MPC.

The following Junos VPN Site Secure features are supported on the MS-MIC and MS-MPC in Release 13.2:

- Dynamic End Points (DEP)
- Encapsulating Security Payload (ESP) protocol
- Dead Peer Detection (DPD) trigger messages
- Sequence Number Rollover notifications
- Static IPsec tunnels with next-hop-style and interface-style service sets

However, in Junos OS, Release 13.2, the Junos VPN Site Secure support on the MS-MIC and MS-MPC is limited to IPv4 traffic. Passive module tunneling is not supported on MS-MICs and MS-MPCs.

This example shows configuration of two routers, Router 1 and Router 2, that have an IPsec VPN tunnel configured between them.

While configuring the routers, note the following points:

- The IP address you configure for **source-address** under the **[edit services ipsec-vpn rule name term term from]** hierarchy level on Router 1 must be the same as the IP address you configure for **destination-address** under the same hierarchy on Router 2, and vice versa.
- The IP address of the **remote-gateway** you configure under the **[edit services ipsec-vpn rule name term term then]** hierarchy level should match the IP address of the **local-gateway** you configure under the **[edit services service-set name ipsec-vpn-options]** hierarchy level of Router 2, and vice versa.

## Configuration

This section contains:

- [Configuring Router 1 on page 478](#)
- [Configuring Router 2 on page 481](#)

### 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 Interfaces on Router 1

```
set interfaces ms-4/0/0 unit 0 family inet
set interfaces ms-4/0/0 unit 1 family inet
set interfaces ms-4/0/0 unit 1 family inet6
```

	<pre> set interfaces ms-4/0/0 unit 1 service-domain inside set interfaces ms-4/0/0 unit 2 family inet set interfaces ms-4/0/0 unit 2 family inet6 set interfaces ms-4/0/0 unit 2 service-domain outside set interfaces xe-0/2/0 unit 0 family inet address 10.0.1/30 </pre>
Configuring IPsec VPN Service on Router 1	<pre> set services ipsec-vpn rule vpn_rule_ms_4_0_01 term term11 from source-address 30.0.0.0/16 set services ipsec-vpn rule vpn_rule_ms_4_0_01 term term11 from destination-address 80.0.0.0/16 set services ipsec-vpn rule vpn_rule_ms_4_0_01 term term11 then remote-gateway 10.0.1.2 set services ipsec-vpn rule vpn_rule_ms_4_0_01 term term11 then dynamic ike-policy ike_policy_ms_4_0_0 set services ipsec-vpn rule vpn_rule_ms_4_0_01 term term11 then dynamic ipsec-policy ipsec_policy_ms_4_0_0 set services ipsec-vpn rule vpn_rule_ms_4_0_01 term term11 then anti-replay-window-size 4096 set services ipsec-vpn rule vpn_rule_ms_4_0_01 match-direction input set services ipsec-vpn ipsec proposal ipsec_proposal_ms_4_0_0 protocol esp set services ipsec-vpn ipsec proposal ipsec_proposal_ms_4_0_0 authentication-algorithm hmac-sha1-96 set services ipsec-vpn ipsec proposal ipsec_proposal_ms_4_0_0 encryption-algorithm 3des-cbc set services ipsec-vpn ipsec policy ipsec_policy_ms_4_0_0 perfect-forward-secrecy keys group2 set services ipsec-vpn ipsec policy ipsec_policy_ms_4_0_0 proposals ipsec_proposal_ms_4_0_0 set services ipsec-vpn ike proposal ike_proposal_ms_4_0_0 authentication-method pre-shared-keys set services ipsec-vpn ike proposal ike_proposal_ms_4_0_0 dh-group group2 set services ipsec-vpn ike policy ike_policy_ms_4_0_0 version 2 set services ipsec-vpn ike policy ike_policy_ms_4_0_0 proposals ike_proposal_ms_4_0_0 set services ipsec-vpn ike policy ike_policy_ms_4_0_0 pre-shared-key ascii-text secret-data </pre>
Configuring a Service Set on Router 1	<pre> set services service-set ipsec_ss_ms_4_0_01 next-hop-service inside-service-interface ms-4/0/0.1 set services service-set ipsec_ss_ms_4_0_01 next-hop-service outside-service-interface ms-4/0/0.2 set services service-set ipsec_ss_ms_4_0_01 ipsec-vpn-options local-gateway 10.0.1.1 set services service-set ipsec_ss_ms_4_0_01 ipsec-vpn-rules vpn_rule_ms_4_0_01 </pre>
Configuring Routing Options on Router 1	<pre> set routing-options static route 80.0.0.0/16 next-hop ms-4/0/0.1 </pre>
Configuring Interfaces on Router 2	<pre> set interfaces ms-1/0/0 unit 0 family inet set interfaces ms-1/0/0 unit 1 family inet set interfaces ms-1/0/0 unit 1 family inet6 set interfaces ms-1/0/0 unit 1 service-domain inside set interfaces ms-1/0/0 unit 2 family inet set interfaces ms-1/0/0 unit 2 family inet6 set interfaces ms-1/0/0 unit 2 service-domain outside set interfaces ge-2/0/0 unit 0 family inet address 10.0.1.2/30 </pre>
Configuring IPsec VPN Service on Router 2	<pre> set services ipsec-vpn rule vpn_rule_ms_5_2_01 term term11 from source-address 80.0.0.0/16 </pre>

```

set services ipsec-vpn rule vpn_rule_ms_5_2_01 term term11 from destination-address
  30.0.0.0/16
set services ipsec-vpn rule vpn_rule_ms_5_2_01 term term11 then remote-gateway 10.0.1.1
set services ipsec-vpn rule vpn_rule_ms_5_2_01 term term11 then dynamic ike-policy
  ike_policy_ms_5_2_0
set services ipsec-vpn rule vpn_rule_ms_5_2_01 term term11 then dynamic ipsec-policy
  ipsec_policy_ms_5_2_0
set services ipsec-vpn rule vpn_rule_ms_5_2_01 term term11 then anti-replay-window-size
  4096
set services ipsec-vpn rule vpn_rule_ms_5_2_01 match-direction input
set services ipsec-vpn ipsec proposal ipsec_proposal_ms_5_2_0 protocol esp
set services ipsec-vpn ipsec proposal ipsec_proposal_ms_5_2_0 authentication-algorithm
  hmac-sha1-96
set services ipsec-vpn ipsec proposal ipsec_proposal_ms_5_2_0 encryption-algorithm
  3des-cbc
set services ipsec-vpn ipsec policy ipsec_policy_ms_5_2_0 perfect-forward-secrecy keys
  group2
set services ipsec-vpn ipsec policy ipsec_policy_ms_5_2_0 proposals
  ipsec_proposal_ms_5_2_0
set services ipsec-vpn ike proposal ike_proposal_ms_5_2_0 authentication-method
  pre-shared-keys
set services ipsec-vpn ike proposal ike_proposal_ms_5_2_0 dh-group group2
set services ipsec-vpn ike policy ike_policy_ms_5_2_0 version 2
set services ipsec-vpn ike policy ike_policy_ms_5_2_0 proposals ike_proposal_ms_5_2_0
set services ipsec-vpn ike policy ike_policy_ms_5_2_0 pre-shared-key ascii-text secret-data
set services ipsec-vpn establish-tunnels immediately

```

#### Configuring a Service Set on Router 2

```

set services service-set ipsec_ss_ms_5_2_01 next-hop-service inside-service-interface
  ms-1/0/0.1
set services service-set ipsec_ss_ms_5_2_01 next-hop-service outside-service-interface
  ms-1/0/0.2
set services service-set ipsec_ss_ms_5_2_01 ipsec-vpn-options local-gateway 10.0.1.2
set services service-set ipsec_ss_ms_5_2_01 ipsec-vpn-rules vpn_rule_ms_5_2_01

```

#### Configuring Routing Options on Router 2

```

set routing-options static route 30.0.0.0/16 next-hop ms-1/0/0.1

```

### Configuring Router 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*.



**NOTE:** Starting with Release 13.2, the Junos OS extension-provider packages come preinstalled on multiservices MICs and MPCs (MS-MICs and MS-MPCs). The adaptive-services configuration at the [edit chassis fpc number pic number] hierarchy level is preconfigured on these cards.

1. Configure the interface properties.

```

user@router1# set interfaces ms-4/0/0 unit 0 family inet
user@router1# set interfaces ms-4/0/0 unit 1 family inet

```



```

user@router1# set interfaces ms-4/0/0 unit 1 family inet6
user@router1# set interfaces ms-4/0/0 unit 1 service-domain inside
user@router1# set interfaces ms-4/0/0 unit 2 family inet
user@router1# set interfaces ms-4/0/0 unit 2 family inet6
user@router1# set interfaces ms-4/0/0 unit 2 service-domain outside
user@router1# set interfaces xe-0/2/0 unit 0 family inet address 10.0.1.1/30

```

2. Configure IPsec properties.

```

user@router1# set services ipsec-vpn rule vpn_rule_ms_4_0_01 term term11 from
source-address 30.0.0.0/16
user@router1# set services ipsec-vpn rule vpn_rule_ms_4_0_01 term term11 from
destination-address 80.0.0.0/16
user@router1# set services ipsec-vpn rule vpn_rule_ms_4_0_01 term term11 then
remote-gateway 10.0.1.2
user@router1# set services ipsec-vpn rule vpn_rule_ms_4_0_01 term term11 then
dynamic ike-policy ike_policy_ms_4_0_0
user@router1# set services ipsec-vpn rule vpn_rule_ms_4_0_01 term term11 then
dynamic ipsec-policy ipsec_policy_ms_4_0_0
user@router1# set services ipsec-vpn rule vpn_rule_ms_4_0_01 term term11 then
anti-replay-window-size 4096
user@router1# set services ipsec-vpn rule vpn_rule_ms_4_0_01 match-direction
input
user@router1# set services ipsec-vpn ipsec proposal ipsec_proposal_ms_4_0_0
protocol esp
user@router1# set services ipsec-vpn ipsec proposal ipsec_proposal_ms_4_0_0
authentication-algorithm hmac-sha1-96
user@router1# set services ipsec-vpn ipsec proposal ipsec_proposal_ms_4_0_0
encryption-algorithm 3des-cbc
user@router1# set services ipsec-vpn ipsec policy ipsec_policy_ms_4_0_0
perfect-forward-secrecy keys group2
user@router1# set services ipsec-vpn ipsec policy ipsec_policy_ms_4_0_0 proposals
ipsec_proposal_ms_4_0_0
user@router1# set services ipsec-vpn ike proposal ike_proposal_ms_4_0_0
authentication-method pre-shared-keys
user@router1# set services ipsec-vpn ike proposal ike_proposal_ms_4_0_0 dh-group
group2
user@router1# set services ipsec-vpn ike policy ike_policy_ms_4_0_0 version 2
user@router1# set services ipsec-vpn ike policy ike_policy_ms_4_0_0 proposals
ike_proposal_ms_4_0_0
user@router1# set services ipsec-vpn ike policy ike_policy_ms_4_0_0 pre-shared-key
ascii-text secret-key

```

3. Configure a service set.

```

user@router1# set services service-set ipsec_ss_ms_4_0_01 next-hop-service
inside-service-interface ms-4/0/0.1
user@router1# set services service-set ipsec_ss_ms_4_0_01 next-hop-service
outside-service-interface ms-4/0/0.2
user@router1# set services service-set ipsec_ss_ms_4_0_01 ipsec-vpn-options
local-gateway 10.0.1.1
user@router1# set services service-set ipsec_ss_ms_4_0_01 ipsec-vpn-rules
vpn_rule_ms_4_0_01

```

4. Configure routing options.

```

user@router1# set routing-options static route 80.0.0.0/16 next-hop ms-4/0/0.1

```

**Results** From the configuration mode of Router 1, confirm your configuration by entering the **show interfaces**, **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@router1# show interfaces
ms-4/0/0 {
  unit 0 {
    family inet;
  }
  unit 1 {
    family inet;
    family inet6;
    service-domain inside;
  }
  unit 2 {
    family inet;
    family inet6;
    service-domain outside;
  }
}
xe-0/2/0 {
  unit 0 {
    family inet {
      address 10.0.1.1/30;
    }
  }
}

user@router1# show services ipsec-vpn
rule vpn_rule_ms_4_0_01 {
  term term11 {
    from {
      source-address {
        30.0.0.0/16;
      }
      destination-address {
        80.0.0.0/16;
      }
    }
    then {
      remote-gateway 10.0.1.2;
      dynamic {
        ike-policy ike_policy_ms_4_0_0;
        ipsec-policy ipsec_policy_ms_4_0_0;
      }
      anti-replay-window-size 4096;
    }
  }
  match-direction input;
}
ipsec {
  proposal ipsec_proposal_ms_4_0_0 {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm 3des-cbc;
```

```

    }
    policy ipsec_policy_ms_4_0_0 {
        perfect-forward-secrecy {
            keys group2;
        }
        proposals ipsec_proposal_ms_4_0_0;
    }
}
ike {
    proposal ike_proposal_ms_4_0_0 {
        authentication-method pre-shared-keys;
        dh-group group2;
    }
    policy ike_policy_ms_4_0_0 {
        version 2;
        proposals ike_proposal_ms_4_0_0;
        pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
    }
}

user@router1# show services service-set
ipsec_ss_ms_4_0_01 {
    next-hop-service {
        inside-service-interface ms-4/0/0.1;
        outside-service-interface ms-4/0/0.2;
    }
    ipsec-vpn-options {
        local-gateway 10.0.1.1;
    }
    ipsec-vpn-rules vpn_rule_ms_4_0_01;
}

```

## Configuring Router 2

### Step-by-Step Procedure

1. Configure the interfaces.
 

```

user@router2# set interfaces ms-1/0/0 services-options inactivity-non-tcp-timeout 600
user@router2# set interfaces ms-1/0/0 unit 0 family inet
user@router2# set interfaces ms-1/0/0 unit 1 family inet
user@router2# set interfaces ms-1/0/0 unit 1 family inet6
user@router2# set interfaces ms-1/0/0 unit 1 service-domain inside
user@router2# set interfaces ms-1/0/0 unit 2 family inet
user@router2# set interfaces ms-1/0/0 unit 2 family inet6
user@router2# set interfaces ms-1/0/0 unit 2 service-domain outside
user@router2# set interfaces ge-2/0/0 unit 0 family inet address 10.0.1.2/30
      
```
2. Configure IPsec properties.
 

```

user@router2# set services ipsec-vpn rule vpn_rule_ms_5_2_01 term term11 from source-address 80.0.0.0/16
user@router2# set services ipsec-vpn rule vpn_rule_ms_5_2_01 term term11 from destination-address 30.0.0.0/16
user@router2# set services ipsec-vpn rule vpn_rule_ms_5_2_01 term term11 then remote-gateway 10.0.1.1
user@router2# set services ipsec-vpn rule vpn_rule_ms_5_2_01 term term11 then dynamic ike-policy ike_policy_ms_5_2_0
      
```

```

user@router2# set services ipsec-vpn rule vpn_rule_ms_5_2_01 term term11 then
dynamic ipsec-policy ipsec_policy_ms_5_2_0
user@router2# set services ipsec-vpn rule vpn_rule_ms_5_2_01 term term11 then
anti-replay-window-size 4096
user@router2# set services ipsec-vpn rule vpn_rule_ms_5_2_01 match-direction
input
user@router2# set services ipsec-vpn ipsec proposal ipsec_proposal_ms_5_2_0
protocol esp
user@router2# set services ipsec-vpn ipsec proposal ipsec_proposal_ms_5_2_0
authentication-algorithm hmac-sha1-96
user@router2# set services ipsec-vpn ipsec proposal ipsec_proposal_ms_5_2_0
encryption-algorithm 3des-cbc
user@router2# set services ipsec-vpn ipsec policy ipsec_policy_ms_5_2_0
perfect-forward-secrecy keys group2
user@router2# set services ipsec-vpn ipsec policy ipsec_policy_ms_5_2_0 proposals
ipsec_proposal_ms_5_2_0
user@router2# set services ipsec-vpn ike proposal ike_proposal_ms_5_2_0
authentication-method pre-shared-keys
user@router2# set services ipsec-vpn ike proposal ike_proposal_ms_5_2_0 dh-group
group2
user@router2# set services ipsec-vpn ike policy ike_policy_ms_5_2_0 version 2
user@router2# set services ipsec-vpn ike policy ike_policy_ms_5_2_0 proposals
ike_proposal_ms_5_2_0
user@router2# set services ipsec-vpn ike policy ike_policy_ms_5_2_0 pre-shared-key
ascii-text "$ABC123"
user@router2# set services ipsec-vpn establish-tunnels immediately

```

3. Configure a service set.

```

user@router2# set services service-set ipsec_ss_ms_5_2_01 next-hop-service
inside-service-interface ms-1/0/0.1
user@router2# set services service-set ipsec_ss_ms_5_2_01 next-hop-service
outside-service-interface ms-1/0/0.2
user@router2# set services service-set ipsec_ss_ms_5_2_01 ipsec-vpn-options
local-gateway 10.0.1.2
user@router2# set services service-set ipsec_ss_ms_5_2_01 ipsec-vpn-rules
vpn_rule_ms_5_2_01

```

4. Configure routing options.

```

user@router2# set routing-options static route 30.0.0.0/16 next-hop ms-1/0/0.1

```

**Results** From the configuration mode of Router 2, confirm your configuration by entering the **show interfaces**, **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@router2# show interfaces
ms-1/0/0 {
  unit 0 {
    family inet;
  }
  unit 1 {
    family inet;
    family inet6;
    service-domain inside;
  }
}

```

```

    }
    unit 2 {
        family inet;
        family inet6;
        service-domain outside;
    }
}
ge-2/0/0 {
    unit 0 {
        family inet {
            address 10.0.1.2/30;
        }
    }
}

user@router2# show services ipsec-vpn
rule vpn_rule_ms_5_2_01 {
    term term11 {
        from {
            source-address {
                80.0.0.0/16;
            }
            destination-address {
                30.0.0.0/16;
            }
        }
        then {
            remote-gateway 10.0.1.1;
            dynamic {
                ike-policy ike_policy_ms_5_2_0;
                ipsec-policy ipsec_policy_ms_5_2_0;
            }
            anti-replay-window-size 4096;
        }
    }
    match-direction input;
}

ipsec {
    proposal ipsec_proposal_ms_5_2_0 {
        protocol esp;
        authentication-algorithm hmac-sha1-96;
        encryption-algorithm 3des-cbc;
    }
    policy ipsec_policy_ms_5_2_0 {
        perfect-forward-secrecy {
            keys group2;
        }
        proposals ipsec_proposal_ms_5_2_0;
    }
}

ike {
    proposal ike_proposal_ms_5_2_0 {
        authentication-method pre-shared-keys;
        dh-group group2;
    }
    policy ike_policy_ms_5_2_0 {

```

```
        version 2;
        proposals ike_proposal_ms_5_2_0;
        pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
    }
}
establish-tunnels immediately;

user@router2# show services service-set
ipsec_ss_ms_5_2_01 {
    next-hop-service {
        inside-service-interface ms-1/0/0.1;
        outside-service-interface ms-1/0/0.2;
    }
    ipsec-vpn-options {
        local-gateway 10.0.1.2;
    }
    ipsec-vpn-rules vpn_rule_ms_5_2_01;
}

user@router2 #show routing-options
static {
    route 30.0.0.0/16 next-hop ms-1/0/0.1;
}
```

## Verification

- [Verifying Tunnel Creation on page 484](#)
- [Verifying Traffic Flow Through the DEP Tunnel on page 485](#)
- [Verifying IPsec Security Associations for the Service Set on page 486](#)

### Verifying Tunnel Creation

---

**Purpose** Verify that Dynamic End Points are created.

**Action** Run the following command on Router 1:

```
user@router1 >show services ipsec-vpn ipsec security-associations detail
Service set: ipsec_ss_ms_4_0_01, IKE Routing-instance: default

Rule: vpn_rule_ms_4_0_01, Term: term11, Tunnel index: 1
Local gateway: 10.0.1.1, Remote gateway: 10.0.1.2
IPSec inside interface: ms-4/0/0.1, Tunnel MTU: 1500
Local identity: ipv4_subnet(any:0,[0..7]=30.0.0.0/16)
Remote identity: ipv4_subnet(any:0,[0..7]=80.0.0.0/16)

Direction: inbound, SPI: 112014862, AUX-SPI: 0
Mode: tunnel, Type: dynamic, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Soft lifetime: Expires in 24556 seconds
Hard lifetime: Expires in 25130 seconds
Anti-replay service: Enabled, Replay window size: 4096

Direction: outbound, SPI: 1469281276, AUX-SPI: 0
Mode: tunnel, Type: dynamic, State: Installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Soft lifetime: Expires in 24556 seconds
Hard lifetime: Expires in 25130 seconds
Anti-replay service: Enabled, Replay window size: 4096
```

**Meaning** The output shows that the IPSec SAs are up on the router with their state as Installed. The IPSec tunnel is up and ready to send traffic over the tunnel.

### Verifying Traffic Flow Through the DEP Tunnel

**Purpose** Verify traffic flow across the newly-created DEP tunnel.

**Action** Run the following command on Router 2:

```
user@router2> show services ipsec-vpn ipsec statistics
PIC: ms-1/0/0, Service set: ipsec_ss_ms_5_2_01

ESP Statistics:
  Encrypted bytes:      153328
  Decrypted bytes:      131424
  Encrypted packets:    2738
  Decrypted packets:    2738
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
```

### Verifying IPsec Security Associations for the Service Set

---

**Purpose** Verify that the security associations configured for the service set are functioning correctly.

**Action** Run the following command on Router 2:

```
user@router2> show services ipsec-vpn ipsec security-associations ipsec_ss_ms_5_2_01
Service set: ipsec_ss_ms_5_2_01, IKE Routing-instance: default
```

```
Rule: vpn_rule_ms_5_2_01, Term: term11, Tunnel index: 1
Local gateway: 10.0.1.2., Remote gateway: 10.0.1.1
IPSec inside interface: ms-1/0/0.1, Tunnel MTU: 1500
  Direction SPI      AUX-SPI      Mode      Type      Protocol
  inbound  1612447024  0          tunnel    dynamic   ESP
  outbound 1824720964  0          tunnel    dynamic   ESP
```



# Enhancing Security with Static IPsec over VRF

- [Example: Configuring Statically Assigned IPsec Tunnels over a VRF Instance on page 487](#)

## 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 487](#)
- [Overview on page 487](#)
- [Configuration on page 487](#)

### 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 488](#)
- [Results on page 490](#)

### 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 "$ABC123"; ## 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;
```

- Related Documentation**
- [Understanding Junos VPN Site Secure on page 401](#)
  - [Configuring Security Associations on page 415](#)
  - [Configuring IPsec Proposals on page 445](#)
  - [Configuring IKE Proposals on page 435](#)





# Dynamically Assigning Tunnels Using Junos VPN Site Secure

- [Configuring Dynamic Endpoints for IPsec Tunnels on page 495](#)
- [Requesting for and Installing a Digital Certificates on Your Router on page 501](#)
- [Example: Configuring Dynamically Assigned Policy Based Tunnels on page 503](#)
- [Example: Configuring IKE Dynamic SAs on page 509](#)
- [Example: IKE Dynamic SA Configuration with Digital Certificates on page 525](#)

## 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 496](#)
- [Implicit Dynamic Rules on page 496](#)
- [Reverse Route Insertion on page 497](#)
- [Configuring an IKE Access Profile on page 497](#)
- [Referencing the IKE Access Profile in a Service Set on page 499](#)
- [Configuring the Interface Identifier on page 499](#)
- [Default IKE and IPsec Proposals on page 500](#)

## 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 21 on page 500](#); if more than one value is shown, the first value is the default.



**NOTE:** RSA certificates are not supported with dynamic endpoint configuration.

**Table 21: Default IKE and IPsec Proposals for Dynamic Negotiations**

Statement Name	Values
<b>Implicit IKE Proposal</b>	
<b>authentication-method</b>	pre-shared keys
<b>dh-group</b>	group1, group2, group5, group14
<b>authentication-algorithm</b>	sha1, md5, sha-256
<b>encryption-algorithm</b>	3des-cbc, des-cbc, aes-128, aes-192, aes-256
<b>lifetime-seconds</b>	3600 seconds
<b>Implicit IPsec Proposal</b>	
<b>protocol</b>	esp, ah, bundle
<b>authentication-algorithm</b>	hmac-sha1-96, hmac-md5-96
<b>encryption-algorithm</b>	3des-cbc, des-cbc, aes-128, aes-192, aes-256
<b>lifetime-seconds</b>	28,800 seconds (8 hours)

### Related Documentation

- [Configuring IKE Policies on page 439](#)
- [Configuring IPsec Rules on page 452](#)
- [Configuring IKE Proposals on page 435](#)
- [Configuring IPsec Proposals on page 445](#)
- [Configuring Security Associations on page 415](#)

## Requesting for and Installing a Digital Certificates on Your Router

A digital certificate is an electronic means for verifying your identity through a trusted third party, known as a certificate authority (CA). Alternatively, you can use a self-signed certificate to attest to your identity. The CA server you use can be owned and operated by an independent CA or by your own organization, in which case you become your own CA. If you use an independent CA, you must contact them for the addresses of their CA and certificate revocation list (CRL) servers (for obtaining certificates and CRLs) and for the information they require when submitting personal certificate requests. When you are your own CA, you determine this information yourself. The Public Key Infrastructure (PKI) provides an infrastructure for digital certificate management.

- [Requesting a Digital Certificate—Manual Process on page 501](#)

### Requesting a Digital Certificate—Manual Process

To obtain digital certificates manually, you must configure a CA profile, generate a private-public key pair, create a local certificate, and load the certificates on the router. After loading the certificates, they can be referenced in your IPsec-VPN configuration.

This procedure shows how you can configure a CA profile:

1. Configure a CA profile:

```
user@R2# set security pki ca-profile entrust ca-identity entrust enrollment url
http://example.com/cgi-bin/pkiclient.exe
```

After you commit this configuration. The configuration on Router 2 must contain the following:

```
[edit]
security {
  pki {
    ca-profile entrust {
      ca-identity entrust;
      enrollment {
        url http://example.com/cgi-bin/pkiclient.exe;
      }
    }
  }
}
```

2. Certificate revocation list (CRL) verification is enabled by default. You can optionally specify the Lightweight Access Directory (LDAP) server where the CA stores the CRL. The certificate typically includes a certificate distribution point (CDP), which contains information about how to retrieve the CRL for the certificate. The router uses this information to download the CRL automatically. In this example, the LDAP URL is specified, which overrides the location provided in the certificate:

```
user@R2# set security pki ca-profile entrust revocation-check crl url
ldap://10.157.90.185/o=juniper,c=uscertificateRevocationListbase
```

After you commit this configuration. The configuration on Router 2 must contain the following:

```
[edit]
security pki ca-profile entrust {
  revocation-check {
    crl {
      url ldap://10.157.90.185/o=juniper,c=uscertificateRevocationListbase;
    }
  }
}
```

- After you configure the CA profile, request a CA certificate from the trusted CA. In this example, the certificate is enrolled online and installed into the router automatically.

```
user@R2> request security pki ca-certificate enroll ca-profile entrust
Received following certificates:
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
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, CN=First Officer
Fingerprint: 46:71:15:34:f0:a6:41:76:65:81:33:4f:68:47:c4:df:78:b8:e3:3f
Do you want to load the above CA certificate ? [yes,no] (no) yes
```



**NOTE:** If you obtain the CA certificate directly from the CA (for example, as an e-mail attachment or website download), you can install it with the `request security pki ca-certificate load` command.

- Next, you must generate a private-public key pair before you can create a local certificate.

```
user@R2> request security pki generate-key-pair certificate-id local-entrust2
Generated key pair local-entrust2, key size 1024 bits
```

When the key pair is available, generate a local certificate request and send it to the CA for processing.

```
user@R2> request security pki generate-certificate-request
certificate-id local-entrust2 domain-name router2.example.com
filename entrust-req2 subject cn=router2.example.com
Generated certificate request
-----BEGIN CERTIFICATE REQUEST-----
MIIBoTCCAQoCAQAwGjEYMBYGA1UEAxMPdHAXLmp1bm1wZXIubmV0MIGFMA0GCSqG
SIb3DQEBAQUAA4GNADCBiQKBgQCiuFk1Qws1Ud+AqN5DDxRs2kVyKEhh9qoVFnz+
Hz4c9vsy3B8E1wTJ1kmIt2cB3yi fB6zePd+6WYpf57Crwre7YqPk iXM31F6z3YjX
H+1BPNbCxNWYvyrnSyVYDbFj8o0Xyqog8ACDfVL2JBWrPNBYy7imq/K9soDBbAs6
5hZqqwIDAQABoEcwRQYJKoZIhvcNAQkOMTgwNjA0BgnVHQ8BAf8EBAMCB4AwJAYD
VRORAQH/BBowGIIWdHAXLmVuZ2xhYi5qdW5pcGVyLm51dDANBgkqhkiG9w0BAQQF
AAOBgQBC2rq1v5S0QXH7LCb/FdqAL8ZM6GoaNs5d6cGwq4bB6a7UQFgtoH406gQ3G
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)
```





**NOTE:** You can request the creation and installation of a local certificate online with the `request security pki local-certificate enroll` command.

5. The trusted CA digitally signs the local certificate and returns it to you. Copy the certificate file into the router and load the certificate.

```
user@R2> request security pki local-certificate load filename /tmp/router2-cert certificate-id
local-entrust2
Local certificate local-entrust2 loaded successfully
```



**NOTE:** The name of the file sent to you by the CA might not match the name of the certificate identifier. However, the `certificate-id` name must always match the name of the key pair you generated for the router.

#### Related Documentation

- [Example: IKE Dynamic SA Configuration with Digital Certificates on page 525](#)

## 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 503](#)
- [Overview and Topology on page 503](#)
- [Configuration on page 504](#)
- [Verification on page 508](#)

### 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 24 on page 504](#).

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 24: 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 SG1 Service-Set on page 505](#)
- [Results on page 506](#)

### 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 SG1 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	<pre> 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 </pre>
Configuring Service Set	<pre> 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 </pre>
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 "$ABC123"; ## 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 401](#)
- [Configuring Security Associations on page 415](#)
- [Configuring IPsec Policies on page 450](#)
- [Configuring IKE Policies on page 439](#)
- [Tracing Junos VPN Site Secure Operations on page 466](#)

## Example: Configuring IKE Dynamic SAs

This example shows how to configure IKE dynamic SAs and contains the following sections.

- [Requirements on page 509](#)
- [Overview and Topology on page 509](#)
- [Configuration on page 510](#)
- [Verification on page 522](#)

### 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.

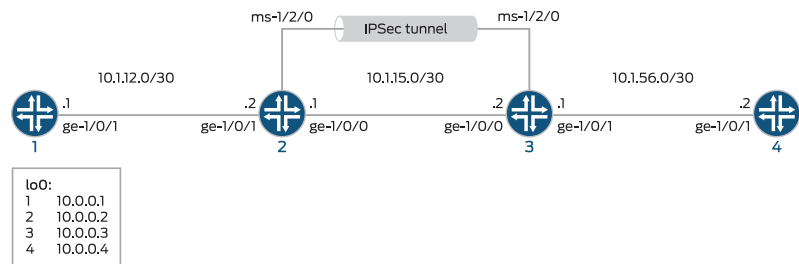
### Overview and Topology

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 25 on page 509](#) 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 25: 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 510](#)
- [Configuring Router 2 on page 511](#)
- [Configuring Router 3 on page 516](#)
- [Configuring Router 4 on page 520](#)

### 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 10.115.2
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 "$ABC123"; ## 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 "$ABC123"; ## 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
```

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-0/0/0
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
```

**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 401](#)
  - [Configuring Security Associations on page 415](#)
  - [Configuring IKE Proposals on page 435](#)
  - [Configuring IKE Policies on page 439](#)
  - [Example: Configuring Manual SAs on page 421](#)

## 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 526](#)
- [Overview on page 526](#)
- [Configuration on page 526](#)
- [Verification on page 540](#)

## 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”](#) on page 501


## 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 26 on page 526](#) 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 26: 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 527](#)
- [Configuring Router 2 on page 528](#)
- [Configuring Router 3 on page 533](#)
- [Configuring Router 4 on page 538](#)

### 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.example.com
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.example.com
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” on page 501](#)

- ```
[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.example.com
user@router2# set ike policy ike-digital-certificates local-certificate local-entrust2
user@router2# set ike policy ike-digital-certificates remote-id fqdn
  router3.example.com
```
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 10.1.15.2
```

```

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
10.1.15.1
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.example.com;
        local-certificate local-entrust2;
        remote-id fqdn router3.example.com;
    }
}
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.example.com
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.example.com
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” on page 501](#)

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” on page 501](#)

```

[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.example.com
user@router3# set ike policy ike-digital-certificates local-certificate local-entrust2
user@router3# set ike policy ike-digital-certificates remote-id fqdn
    router3.example.com

```

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.example.com;
    local-certificate local-entrust3;
    remote-id fqdn router2.example.com;
  }
}
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.example.com, Issued by: juniper
Alternate subject: router3.example.com
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.example.com, Issued by: juniper
Alternate subject: router2.example.com
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.example.com
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.example.com, 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.example.com, Issued by: juniper
Alternate subject: router3.example.com
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.example.com, Issued by: juniper
Alternate subject: router2.example.com
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.example.com
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.example.com, 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 401](#)
- [Configuring Security Associations on page 415](#)
- [Configuring IKE Proposals on page 435](#)

- [Configuring IKE Policies on page 439](#)
- [Example: Configuring IKE Dynamic SAs on page 509](#)
- [Example: Configuring Manual SAs on page 421](#)
- [Requesting for and Installing a Digital Certificates on Your Router on page 501](#)

# Enabling IPsec for the Services SDK

- [Configuring Junos VPN Site Secure or IPSec VPN on page 549](#)

## 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 415](#)
  - [Service Sets for IPsec Tunnels on page 460](#)





## PART 7

# Alleviating Congestion and Controlling Service Using CoS

- [Class of Service Overview on page 553](#)
- [Class of Service Configuration Overview on page 555](#)
- [Configuring Class of Service on LSQ Interfaces on page 565](#)



# Class of Service Overview

- [Class of Service Overview on page 553](#)

## 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 555](#)
- [Configuring CoS Rules on page 556](#)
- [Configuring CoS Rule Sets on page 561](#)
- [Examples: Configuring CoS on Services Interfaces on page 561](#)



# Class of Service Configuration Overview

- [Restrictions and Cautions for CoS Configuration on Services Interfaces on page 555](#)
- [Configuring CoS Rules on page 556](#)
- [Configuring CoS Rule Sets on page 561](#)
- [Examples: Configuring CoS on Services Interfaces on page 561](#)

## 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 553](#)
- [Configuring CoS Rules on page 556](#)

- [Configuring CoS Rule Sets on page 561](#)
- [Examples: Configuring CoS on Services Interfaces on page 561](#)

## 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 557](#)
- [Configuring Match Conditions In CoS Rules on page 557](#)
- [Configuring Actions in CoS Rules on page 558](#)
- [Example: Configuring CoS Rules on page 560](#)

## 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 9](#).

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 363](#).

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 325](#).

- 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 561 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 559](#)
- [Configuring Reflexive and Reverse CoS Rule Actions on page 560](#)

### 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 553](#)
- [Restrictions and Cautions for CoS Configuration on Services Interfaces on page 555](#)
- [Configuring CoS Rule Sets on page 561](#)
- [Examples: Configuring CoS on Services Interfaces on page 561](#)

## 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 553](#)
- [Restrictions and Cautions for CoS Configuration on Services Interfaces on page 555](#)
- [Configuring CoS Rules on page 556](#)
- [Examples: Configuring CoS on Services Interfaces on page 561](#)

## 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.

<b>Mapping Forwarding-Class Name to Forwarding-Class ID</b>	<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>
<b>Mapping Forwarding-Class Name to Queue Number</b>	<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>
<b>Mapping Diffserv Code Point Aliases to DSCP Bits</b>	<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 553](#)
- [Restrictions and Cautions for CoS Configuration on Services Interfaces on page 555](#)
- [Configuring CoS Rules on page 556](#)
- [Configuring CoS Rule Sets on page 561](#)



# Configuring Class of Service on LSQ Interfaces

- [Link Services Configuration for Junos Interfaces on page 565](#)
- [Configuring CoS Scheduling Queues on Logical LSQ Interfaces on page 566](#)
- [Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 570](#)
- [Configuring Link Services and CoS on Services PICs on page 572](#)
- [Oversubscribing Interface Bandwidth on LSQ Interfaces on page 575](#)
- [Configuring Guaranteed Minimum Rate on LSQ Interfaces on page 580](#)

## Link Services Configuration for Junos Interfaces

---

This topic provides links to topics explaining link services configuration for the following interface types:

- For information about configuring LSQ interface redundancy across multiple routers using SONET APS interfaces, see [“Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS” on page 589](#)
- For information about configuring LSQ interface redundancy in a single router using SONET APS interfaces, see [“Configuring LSQ Interface Redundancy in a Single Router Using SONET APS” on page 592](#)
- For information about configuring LSQ interface redundancy in a single router using Virtual Interfaces, see [“Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces” on page 592](#)
- For information about configuring CoS scheduling queues on Logical LSQ interfaces, see [“Configuring CoS Scheduling Queues on Logical LSQ Interfaces” on page 566](#)
- For information about configuring CoS fragmentation by forwarding class on LSQ interfaces, see [“Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces” on page 570](#)
- For information about reserving bundle bandwidth for Link-Layer overhead on LSQ interfaces, see [“Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces” on page 605](#)

- For information about configuring multiclass MLPPP on LSQ interfaces, see [“Configuring Multiclass MLPPP on LSQ Interfaces” on page 606](#)
- For information about oversubscribing interface bandwidth on LSQ interfaces, see [“Oversubscribing Interface Bandwidth on LSQ Interfaces” on page 575](#)
- For information about configuring guaranteed minimum rate on LSQ interfaces, see [“Configuring Guaranteed Minimum Rate on LSQ Interfaces” on page 580](#)
- For information about configuring link services and CoS on services PICs, see [“Configuring Link Services and CoS on Services PICs” on page 572](#)
- For information about configuring LSQ interfaces as NxT1 or NxEl bundles using MLPPP, see [“Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using MLPPP” on page 609](#)
- For information about configuring LSQ interfaces as NxT1 or NxEl bundles using FRF.16, see [“Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.16” on page 615](#)
- For information about configuring LSQ interfaces for single fractional T1 or El interfaces using MLPPP and LFI, see [“Configuring LSQ Interfaces for Single Fractional T1 or El Interfaces Using MLPPP and LFI” on page 621](#)
- For information about configuring LSQ interfaces for single fractional T1 or El interfaces using FRF.12, see [“Configuring LSQ Interfaces for Single Fractional T1 or El Interfaces Using FRF.12” on page 626](#)
- For information about configuring LSQ interfaces as NxT1 or NxEl bundles using FRF.15, see [“Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.15” on page 620](#)
- For information about configuring LSQ interfaces for T3 links configured for compressed RTP over MLPPP, see [“Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP” on page 633](#)
- For information about configuring LSQ interfaces as T3 or OC3 bundles using FRF.12, see [“Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12” on page 635](#)
- For information about configuring LSQ interfaces for ATM2 IQ interfaces using MLPPP, see [“Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP” on page 637](#)

**Related  
Documentation**

- [Layer 2 Service Package Capabilities and Interfaces on page 587](#)

---

## 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 568](#).
- **priority**—The transmit priority (low, high, strict-high); for more information, see [“Configuring Scheduler Priority” on page 568](#).
- **shaping-rate**—The subscribed transmit rate; for more information, see [“Configuring Scheduler Shaping Rate” on page 568](#).
- **drop-profile-map**—The random early detection (RED) drop profile; for more information, see [“Configuring Drop Profiles” on page 569](#).

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 618](#). 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 587](#)
- [Configuring Link Services and CoS on Services PICs on page 572](#)
- [Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 570](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

## 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*.

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 606.

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 609](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.16 on page 615](#)
- [Configuring LSQ Interfaces for Single Fractional T1 or El Interfaces Using MLPPP and LFI on page 621](#)

- [Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12 on page 626](#)
- [Configuring LSQ Interfaces as NxT1 or NxT1 Bundles Using FRF.15 on page 620](#)
- [Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 633](#)
- [Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12 on page 635](#)
- [Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP on page 637](#)

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*.

**Related Documentation**

- [Layer 2 Service Package Capabilities and Interfaces on page 587](#)
- [Configuring Link Services and CoS on Services PICs on page 572](#)
- [Configuring CoS Scheduling Queues on Logical LSQ Interfaces on page 566](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

---

## 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* and [“Layer 2 Service Package Capabilities and Interfaces” on page 587](#).

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 587](#)
- [Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS on page 589](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using SONET APS on page 592](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 592](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

## 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 580](#).

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 578](#).

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:

(interface speed) – (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** Apply a traffic-control profile to a logical interface representing a DLCI on an FRF.16 bundle.

```
interfaces {
```

**Scheduling Based on the Logical Interface**

```

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 {
      dlci 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 587](#)
- [Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces on page 605](#)
- [Configuring Guaranteed Minimum Rate on LSQ Interfaces on page 580](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

---

## 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 583](#).

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:

delay-buffer-rate < [8 x 64 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 {
  t1-3/0/1 {
    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 t1-3/0/1 {
  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 587](#)
- [Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces on page 605](#)
- [Oversubscribing Interface Bandwidth on LSQ Interfaces on page 575](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)



## PART 8

# Configuring Interface Redundancy and Bundling on LSQ Interfaces

- [Overview on page 587](#)
- [Configuring Interface Redundancy with SONET APS and Virtual Interfaces on page 589](#)
- [Enabling Bundling on LSQ Interfaces on page 603](#)



# Overview

- [Layer 2 Service Package Capabilities and Interfaces on page 587](#)

## Layer 2 Service Package Capabilities and Interfaces

---

As described in *Enabling Service Packages*, 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 566 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 (CRTP) 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 CRTP 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 668 and “[Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces](#)” on page 570.

- 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 566.



**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 38

# Configuring Interface Redundancy with SONET APS and Virtual Interfaces

- [Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS on page 589](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using SONET APS on page 592](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 592](#)

### 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 590](#)
- [Configuring SONET APS Interoperability with Cisco Systems FRF.16 on page 591](#)
- [Restrictions on APS Redundancy for LSQ Interfaces on page 591](#)

## 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 587](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using SONET APS on page 592](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 592](#)

- [Configuring Link Services and CoS on Services PICs on page 572](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

---

## 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 589. 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 587](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 592](#)
- [Configuring Link Services and CoS on Services PICs on page 572](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

---

## 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 593](#)
- [Restrictions on Redundant LSQ Interfaces on page 594](#)
- [Configuring Link State Replication for Redundant Link PICs on page 595](#)
- [Examples: Configuring Redundant LSQ Interfaces for Failure Recovery on page 597](#)

## 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 589](#). 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 600](#).



**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 587](#)
- [Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS on page 589](#)
- [Configuring LSQ Interface Redundancy in a Single Router Using SONET APS on page 592](#)
- [Configuring Link Services and CoS on Services PICs on page 572](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)



# Enabling Bundling on LSQ Interfaces

- [Inline MLPPP for WAN Interfaces Overview on page 603](#)
- [Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces on page 605](#)
- [Configuring Multiclass MLPPP on LSQ Interfaces on page 606](#)
- [Enabling Inline LSQ Services on page 607](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using MLPPP on page 609](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.16 on page 615](#)
- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.15 on page 620](#)
- [Configuring LSQ Interfaces for Single Fractional T1 or El Interfaces Using MLPPP and LFI on page 621](#)
- [Configuring LSQ Interfaces for Single Fractional T1 or El Interfaces Using FRF.12 on page 626](#)
- [Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 633](#)
- [Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12 on page 635](#)
- [Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP on page 637](#)

## 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.

Starting in Junos OS Release 15.1, you can configure inline MLPPP interfaces on MX80, MX104, MX240, MX480, and MX960 routers with Channelized E1/T1 Circuit Emulation MICs. A maximum of up to eight inline MLPPP interface bundles are supported on Channelized E1/T1 Circuit Emulation MICs, similar to the support for inline MLPPP bundles on other MICs with which they are compatible.

Configuring inline MLPPP for WAN interfaces benefits the following services:

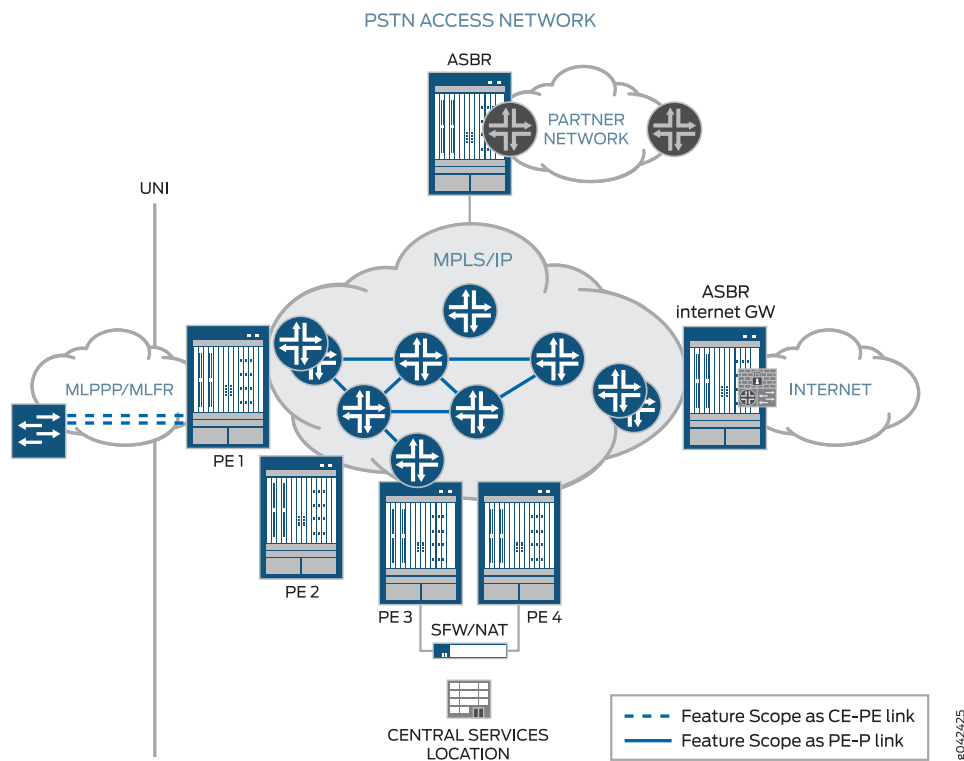
- CE-PE link for Layer 3 VPN and DIA service with public switched telephone networks (PSTN)-based access networks.
- PE-P link when PSTN is used for MPLS networks.

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 27: 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*
- *Example: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces*

#### Related Documentation

- *Configuring the Junos OS to Support the Link Services PIC*
- [Enabling Inline LSQ Services on page 607](#)
- *Enabling MLPPP Link Fragmentation and Interleaving*
- *Example: Configuring Multilink Frame Relay FRF.15*
- *Example: Configuring Multilink Frame Relay FRF.16*
- *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 587](#)
- [Oversubscribing Interface Bandwidth on LSQ Interfaces on page 575](#)
- [Configuring Guaranteed Minimum Rate on LSQ Interfaces on page 580](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

---

## 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*.

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 666](#).

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 587](#)
- [Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on page 609](#)
- [Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP on page 637](#)
- [Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 633](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

## 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/0`. 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/0: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`, and `lsq-1/0/0:0` 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 603](#)
- [Link Services IQ Interfaces](#)
- [Link and Multilink Services Interfaces Feature Guide for Routing Devices](#)
- [mlfr-uni-nni-bundles-inline on page 829](#)
- [multi-link-layer-2-inline on page 831](#)

## Configuring LSQ Interfaces as NxT1 or Nx E1 Bundles Using MLPPP

To configure an  $N \times T1$  bundle using MLPPP, you aggregate  $N$  different T1 links into a bundle. The  $N \times T1$  bundle is called a logical interface, because it can represent, for example, a routing adjacency. To aggregate T1 links into 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 612](#).



**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 606](#). For more information about fragmentation maps, see [“Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces” on page 570](#).

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*.

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 587](#)
  - [Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16 on page 615](#)



- [Configuring LSQ Interfaces as NxT1 or NxEl Bundles Using FRF.15 on page 620](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

## Configuring LSQ Interfaces as NxT1 or NxEl 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 {
    dlci dlci-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 618](#).

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/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*.

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 587](#)
- [Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on page 609](#)
- [Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.15 on page 620](#)
- [Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 633](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

---

## Configuring LSQ Interfaces as NxT1 or NxE1 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 E1 Interfaces Using FRF.12” on page 626. 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 587](#)
- [Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on page 609](#)
- [Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16 on page 615](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

## 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 624.



**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 570](#).

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/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*.

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 587](#)
- [Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12 on page 626](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

## 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 629](#).



**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 570](#).

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*.

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 00.00.5e.00.53.56;
      }
    }
  }
}
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 00.00.5e.00.53.56;
      }
    }
  }
  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 587](#)
- [Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI on page 621](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

## 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 666](#).

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 587](#)
- [Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on page 609](#)
- [Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI on page 621](#)
- [Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP on page 637](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)

## 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 606](#).

```

[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 626](#).

#### Related Documentation

- [Layer 2 Service Package Capabilities and Interfaces on page 587](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)
- [Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 633](#)

## 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 587](#)
- [Link Services Configuration for Junos Interfaces on page 565](#)



## PART 9

# Enabling Load Balancing and High Availability Using Multiservices Interfaces

- [Enabling Load Balancing and High Availability Using Multiservices Interfaces on page 643](#)



## CHAPTER 40

# Enabling Load Balancing and High Availability Using Multiservices Interfaces

- [Understanding Aggregated Multiservices Interfaces on page 643](#)
- [Configuring Load Balancing on AMS Infrastructure on page 649](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 652](#)
- [Example: Configuring Next-Hop Style Services on an Aggregated Multiservices Interface on page 657](#)
- [Example: Configuring Static Source Translation on AMS Infrastructure on page 660](#)

## Understanding Aggregated Multiservices Interfaces

---

This topic contains the following sections:

- [Aggregated Multiservices Interface on page 643](#)
- [IPv6 Traffic on AMS Interfaces Overview on page 647](#)
- [Member Failure Options and High Availability Settings on page 648](#)

### 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**.

When you want **services-options** settings to apply uniformly to all members, configure **services-options** at the **ams** interface level. If you need different settings for individual members (for example, because of a syslog configuration), configure **services-options** at the member-interface level.



**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:** 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:** 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 Nth 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 NAPT44).

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/O]** 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 652](#)
- [Example: Configuring Next-Hop Style Services on an Aggregated Multiservices Interface on page 657](#)

---

## 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 660](#) for more details on configuring NAT flows for load balancing.

### Related Documentation

- [Understanding Aggregated Multiservices Interfaces on page 643](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 652](#)
- [Example: Configuring Next-Hop Style Services on an Aggregated Multiservices Interface on page 657](#)

- [Example: Configuring Static Source Translation on AMS Infrastructure on page 660](#)

## Example: Configuring an Aggregated Multiservices Interface (AMS)

---

- [Hardware and Software Requirements on page 652](#)
- [Overview on page 652](#)
- [Configuration on page 653](#)
- [Verification on page 656](#)

### 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 643](#).

## Configuration

<b>CLI Quick Configuration</b>	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.
<b>Adding Member Interfaces</b>	<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>
<b>Configuring Logical Units</b>	<pre>set interfaces ams0 unit 1 family inet</pre>
<b>Configuring Member Failure Options</b>	<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>
<b>Configuring High Availability Options</b>	<pre>set interfaces ams0 load-balancing-options high-availability-options many-to-one preferred-backup mams-1/0/0</pre>
<b>Configuring Service Set and Interface Services</b>	<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>
<b>Step-by-Step Procedure</b>	<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"> <li>1. Create an aggregated multiservices interface and add member interfaces.</li> </ol>



**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 22: Key Configuration Statements Used in this Example

Statement	Description
<b>member-interface</b>	Adds a member interface (mams) to the AMS bundle.
<b>drop-member-traffic</b>	Specifies that all traffic to a member be dropped in case the member interface fails.
<b>rejoin-timeout</b>	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.
<b>enable-rejoin</b>	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.
<b>preferred-backup</b>	Designates a member interface as the floating backup.
<b>interface-services</b>	Specifies a service interface, an AMS interface in this example, to handle interface services.
<b>hash-keys</b>	Specifies the load-balancing hash keys. You can configure the following hash key values: <b>source-ip</b> , <b>destination-ip</b> , <b>iif</b> (incoming interface), <b>oif</b> (outgoing interface), and <b>protocol</b> .  <b>NOTE:</b> 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 656](#)

### 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 657](#)
  - [Understanding Aggregated Multiservices Interfaces on page 643](#)

## Example: Configuring Next-Hop Style Services on an Aggregated Multiservices Interface

- [Hardware and Software Requirements on page 657](#)
- [Overview on page 657](#)
- [Configuration on page 657](#)

### 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 652](#)
  - [Understanding Aggregated Multiservices Interfaces on page 643](#)

---

## 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 649](#)
- [Understanding Aggregated Multiservices Interfaces on page 643](#)



## PART 10

# Handling VoIP, HTTP, and Layer 2 Traffic

- [Handling VoIP Traffic Using Voice Services on page 665](#)
- [Tunneling PPP Packets Across a Network Using Layer 2 Tunneling on page 675](#)



# Handling VoIP Traffic Using Voice Services

- [Voice Services Overview on page 665](#)
- [Configuring Services Interfaces for Voice Services on page 666](#)
- [Configuring Encapsulation for Voice Services on page 669](#)
- [Configuring Network Interfaces for Voice Services on page 670](#)
- [Examples: Configuring Voice Services on page 671](#)

## 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 666](#).

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 572](#).

### Related Documentation

- [Configuring Services Interfaces for Voice Services on page 666](#)
- [Configuring Encapsulation for Voice Services on page 669](#)

- [Configuring Network Interfaces for Voice Services on page 670](#)
- [Examples: Configuring Voice Services on page 671](#)

## 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 666](#)
- [Configuring Compression of Voice Traffic on page 667](#)
- [Configuring Delay-Sensitive Packet Interleaving on page 668](#)
- [Example: Configuring Compression of Voice Traffic on page 668](#)

### 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 interoperation 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 665](#)

- [Configuring Encapsulation for Voice Services on page 669](#)
- [Configuring Network Interfaces for Voice Services on page 670](#)
- [Examples: Configuring Voice Services on page 671](#)

## 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 671.



**NOTE:** The only protocol type supported with **frame-relay-ppp** encapsulation is **family mlppp**.

### Related Documentation

- [Voice Services Overview on page 665](#)
- [Configuring Services Interfaces for Voice Services on page 666](#)
- [Configuring Network Interfaces for Voice Services on page 670](#)
- [Examples: Configuring Voice Services on page 671](#)

## 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 670](#)
- [Configuring the Compression Interface with PPP Encapsulation on page 670](#)

### 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*.

### 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 665](#)
- [Configuring Services Interfaces for Voice Services on page 666](#)
- [Configuring Encapsulation for Voice Services on page 669](#)
- [Examples: Configuring Voice Services on page 671](#)

## 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 665](#)
- [Configuring Services Interfaces for Voice Services on page 666](#)
- [Configuring Encapsulation for Voice Services on page 669](#)
- [Configuring Network Interfaces for Voice Services on page 670](#)



# Tunneling PPP Packets Across a Network Using Layer 2 Tunneling

- [Layer 2 Tunneling Protocol Overview on page 675](#)
- [L2TP Services Configuration Overview on page 676](#)
- [L2TP Minimum Configuration on page 677](#)
- [Configuring L2TP Tunnel Groups on page 679](#)
- [Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 684](#)
- [AS PIC Redundancy for L2TP Services on page 686](#)
- [Examples: Configuring L2TP Services on page 686](#)
- [Tracing L2TP Operations on page 689](#)

## 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 676.

- Related Documentation**
- [L2TP Services Configuration Overview on page 676](#)
  - [AS PIC Redundancy for L2TP Services on page 686](#)
  - [L2TP Minimum Configuration on page 677](#)
  - [Examples: Configuring L2TP Services on page 686](#)

---

## 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 more information about configuring properties at the **[edit access]** hierarchy level, see the *Junos OS Administration Library for Routing Devices*. For information about L2TP LAC and LNS configurations on MX Series routers for subscriber access, see *L2TP for Subscriber Access Overview* in the *Junos Subscriber Access Configuration Guide*.

---

## 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 "$ABC123"; # SECRET-DATA
    }
    user-group-profile westcoast_users;
  }
}
profile westcoast_bldg_1 {
  authentication-order radius;
}

```



```

}
radius-server {
  192.168.65.63 {
    port 1812;
    secret "$ABC123"; # SECRET-DATA
  }
}

```

**Related  
Documentation**

- [L2TP Services Configuration Overview on page 676](#)
- [Configuring L2TP Tunnel Groups on page 679](#)
- [Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 684](#)
- [Tracing L2TP Operations on page 689](#)
- [Examples: Configuring L2TP Services on page 686](#)

## 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 680](#)
- [Configuring the Local Gateway Address and PIC on page 680](#)
- [Configuring Window Size for L2TP Tunnels on page 681](#)
- [Configuring Timers for L2TP Tunnels on page 681](#)
- [Hiding Attribute-Value Pairs for L2TP Tunnels on page 682](#)
- [Configuring System Logging of L2TP Tunnel Activity on page 682](#)

## 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 686.

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 23 on page 682 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 23: System Log Message Severity Levels**

Severity Level	Description
<b>any</b>	Includes all severity levels
<b>emergency</b>	System panic or other condition that causes the router to stop functioning

Table 23: System Log Message Severity Levels (*continued*)

Severity Level	Description
<b>alert</b>	Conditions that require immediate correction, such as a corrupted system database
<b>critical</b>	Critical conditions, such as hard drive errors
<b>error</b>	Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels
<b>warning</b>	Conditions that warrant monitoring
<b>notice</b>	Conditions that are not errors but might warrant special handling
<b>info</b>	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 [System Log Explorer](#).

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 676](#)
- [L2TP Minimum Configuration on page 677](#)
- [Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 684](#)
- [Tracing L2TP Operations on page 689](#)
- [Examples: Configuring L2TP Services on page 686](#)

## 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 "$ABC123"; # SECRET-DATA
      }
    }
  }
  profile u {
    authentication-order radius;
  }
  radius-server {
    192.168.65.63 {
      port 1812;
      secret "$ABC123"; # 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 676](#)
- [L2TP Minimum Configuration on page 677](#)
- [Configuring L2TP Tunnel Groups on page 679](#)
- [Tracing L2TP Operations on page 689](#)
- [Examples: Configuring L2TP Services on page 686](#)

## 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 20](#). For an example configuration, see [“Examples: Configuring L2TP Services” on page 686](#). For information on operational mode commands, see the [CLI Explorer](#).

### Related Documentation

- [Layer 2 Tunneling Protocol Overview on page 675](#)
- [L2TP Services Configuration Overview on page 676](#)
- [Configuring AS or Multiservices PIC Redundancy on page 20](#)
- [L2TP Minimum Configuration on page 677](#)
- [Examples: Configuring L2TP Services on page 686](#)

## 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 "$ABC123"; # SECRET-DATA
        ppp {
            idle-timeout 22;
            primary-dns 192.168.65.1;
            framed-ip-address 10.12.12.12/32;
            interface-id east;
        }
        group-profile sunnyvale_users;
    }
    client blue {
        chap-secret "$ABC123"; # SECRET-DATA
        group-profile sunnyvale_users;
    }
    authentication-order password;
}
profile sunnyvale_bldg_1_tunnel {
    client test {
        l2tp {
            shared-secret "$ABC123"; # SECRET-DATA
            maximum-sessions-per-tunnel 75;
            interface-id west_shared;
            ppp-authentication chap;
        }
    }
}
```

```
    }
    group-profile sunnyvale_tunnel;
  }
  client production {
    l2tp {
      shared-secret "$ABC123";
      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]
unit0 {
  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 676](#)
- [L2TP Minimum Configuration on page 677](#)
- [Configuring L2TP Tunnel Groups on page 679](#)
- [Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 684](#)
- [Tracing L2TP Operations on page 689](#)

## 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 676](#)
- [L2TP Minimum Configuration on page 677](#)
- [Configuring L2TP Tunnel Groups on page 679](#)
- [Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 684](#)
- [Examples: Configuring L2TP Services on page 686](#)



## PART 11

# Configuration Statements and Operational Commands

- Configuration Statements on page 695
- Operational Commands on page 959





# Configuration Statements

- [\[edit services application-identification\] Hierarchy Level on page 704](#)
- [IPsec Hierarchy Level on page 706](#)
- [adaptive-services-pics on page 709](#)
- [address \(Interfaces\) on page 710](#)
- [address \(Services NAT Pool\) on page 710](#)
- [address-allocation on page 711](#)
- [address-range on page 711](#)
- [aggregation on page 712](#)
- [allow-ip-options on page 713](#)
- [allow-multicast on page 714](#)
- [allow-overlapping-nat-pools on page 714](#)
- [anti-replay-window-size \(Services IPsec VPN\) on page 714](#)
- [anti-replay-window-size \(Services Service Set\) on page 715](#)
- [app-mapping-timeout on page 716](#)
- [application on page 717](#)
- [application-protocol on page 718](#)
- [application-profile on page 720](#)
- [application-set on page 721](#)
- [application-sets \(Services CoS\) on page 721](#)
- [application-sets \(Services IDS\) on page 722](#)
- [application-sets \(Services NAT\) on page 722](#)
- [application-sets \(Services Stateful Firewall\) on page 723](#)
- [applications \(Services ALGs\) on page 723](#)
- [applications \(Services CoS\) on page 724](#)
- [applications \(Services IDS\) on page 724](#)
- [applications \(Services NAT\) on page 725](#)
- [applications \(Services Stateful Firewall\) on page 725](#)
- [authentication on page 726](#)

- [authentication-algorithm \(Services IKE\)](#) on page 727
- [authentication-algorithm \(Services IPsec\)](#) on page 728
- [authentication-method](#) on page 729
- [auxiliary-spi](#) on page 730
- [backup-remote-gateway](#) on page 730
- [bundle](#) on page 731
- [by-destination](#) on page 731
- [by-pair](#) on page 732
- [by-source](#) on page 733
- [bypass-traffic-on-exceeding-flow-limits](#) on page 733
- [bypass-traffic-on-pic-failure](#) on page 734
- [cgn-pic](#) on page 734
- [cisco-interoperability](#) on page 735
- [class](#) on page 736
- [clear-dont-fragment-bit \(Interfaces GRE Tunnels\)](#) on page 737
- [clear-dont-fragment-bit](#) on page 737
- [clear-dont-fragment-bit \(Services NAT Options\)](#) on page 738
- [clear-dont-fragment-bit \(Services Service Set\)](#) on page 738
- [clear-ike-sas-on-pic-restart](#) on page 739
- [clear-ipsec-sas-on-pic-restart](#) on page 739
- [compression](#) on page 740
- [compression-device \(Interfaces\)](#) on page 740
- [copy-dont-fragment-bit \(Services IPsec VPN\)](#) on page 741
- [copy-dont-fragment-bit \(Services Set\)](#) on page 741
- [data \(FTP\)](#) on page 742
- [dead-peer-detection \(Services IPsec VPN\)](#) on page 742
- [description \(Services IPsec VPN\)](#) on page 743
- [destination-address \(Services CoS\)](#) on page 743
- [destination-address \(Services IDS\)](#) on page 744
- [destination-address](#) on page 744
- [destination-address \(Services NAT\)](#) on page 745
- [destination-address \(Services Stateful Firewall\)](#) on page 745
- [destination-address-range \(Services IDS\)](#) on page 746
- [destination-address-range \(Services NAT\)](#) on page 747
- [destination-address-range \(Services Stateful Firewall\)](#) on page 748
- [destination-pool](#) on page 748
- [destination-port](#) on page 749

- [destination-port range on page 750](#)
- [destination-prefix \(Services IDS\) on page 750](#)
- [destination-prefix \(Services NAT\) on page 751](#)
- [destination-prefix-ipv6 on page 751](#)
- [destination-prefix-list \(Services CoS\) on page 752](#)
- [destination-prefix-list \(Services IDS\) on page 752](#)
- [destination-prefix-list \(Services NAT\) on page 753](#)
- [destination-prefix-list \(Services Stateful Firewall\) on page 753](#)
- [destined-port on page 754](#)
- [deterministic-port-block-allocation on page 755](#)
- [dh-group on page 756](#)
- [dial-options on page 757](#)
- [direction on page 758](#)
- [dns-alg-pool on page 758](#)
- [dns-alg-prefix on page 759](#)
- [drop-member-traffic \(Aggregated Multiservices\) on page 759](#)
- [ds-lite on page 760](#)
- [dscp on page 761](#)
- [dynamic on page 761](#)
- [ecmp-alb on page 762](#)
- [ei-mapping-timeout on page 763](#)
- [eif-flow-limit on page 763](#)
- [enable-change-on-ams-redistribution on page 764](#)
- [enable-rejoin \(aggregated Multiservices\) on page 765](#)
- [encapsulation on page 766](#)
- [encryption on page 767](#)
- [encryption-algorithm on page 768](#)
- [establish-tunnels on page 769](#)
- [f-max-period on page 769](#)
- [facility-override \(Service Sets\) on page 770](#)
- [facility-override \(System Log Reporting\) on page 771](#)
- [family \(Aggregated Multiservices\) on page 771](#)
- [family \(Interfaces\) on page 772](#)
- [family \(Voice Services\) on page 773](#)
- [force-entry on page 774](#)
- [forwarding-class \(Services CoS\) on page 774](#)
- [forwarding-class \(Services CoS Fragmentation Properties\) on page 775](#)

- [fragment-limit on page 775](#)
- [fragment-threshold \(Forwarding Class Maps\) on page 776](#)
- [fragment-threshold \(Interfaces LSQ\) on page 777](#)
- [fragmentation-map on page 777](#)
- [fragmentation-maps on page 778](#)
- [from \(Services CoS\) on page 779](#)
- [from \(Services IDS\) on page 780](#)
- [from on page 781](#)
- [from \(Services HCM\) on page 781](#)
- [from \(Services NAT\) on page 782](#)
- [from \(Services Stateful Firewall\) on page 783](#)
- [ftp \(Services CoS\) on page 784](#)
- [hash-keys \(Aggregated Multiservices\) on page 785](#)
- [header-integrity-check on page 788](#)
- [hello-interval on page 789](#)
- [hide-avps on page 790](#)
- [high-availability-options \(aggregated Multiservices\) on page 791](#)
- [hint on page 792](#)
- [host \(L2TP\) on page 792](#)
- [host \(service-set\) on page 793](#)
- [host \(Services HCM\) on page 794](#)
- [hot-standby on page 794](#)
- [icmp-code on page 795](#)
- [icmp-type on page 795](#)
- [ids-rules on page 796](#)
- [ignore-entry on page 796](#)
- [ike on page 797](#)
- [ike-access-profile on page 798](#)
- [inactivity-timeout on page 798](#)
- [initiate-dead-peer-detection on page 799](#)
- [input \(Interfaces\) on page 799](#)
- [interface on page 800](#)
- [interface-service on page 800](#)
- [interfaces \(Aggregated Multiservices\) on page 801](#)
- [interfaces \(Voice Services\) on page 802](#)
- [interval on page 802](#)
- [ipsec on page 803](#)

- [ipsec-inside-interface](#) on page 803
- [ipsec-vpn-options](#) on page 804
- [ipsec-vpn-rules](#) on page 804
- [ipv6-multicast-interfaces](#) on page 805
- [l2tp-access-profile](#) on page 805
- [land-attack-check](#) on page 806
- [learn-sip-register](#) on page 806
- [lifetime-seconds](#) on page 807
- [link-layer-overhead](#) on page 807
- [load-balance](#) on page 808
- [load-balancing-options \(Aggregated Multiservices\)](#) on page 809
- [local-certificate](#) on page 810
- [local-gateway \(IPSec\)](#) on page 811
- [local-gateway \(L2TP LNS\)](#) on page 811
- [local-id](#) on page 812
- [log-prefix \(L2TP\)](#) on page 812
- [log-prefix \(Services\)](#) on page 813
- [logging \(Services\)](#) on page 813
- [logging \(Services IDS\)](#) on page 814
- [lsq-failure-options](#) on page 814
- [manual](#) on page 815
- [many-to-one \(Aggregated Multiservices\)](#) on page 816
- [mapping-refresh](#) on page 817
- [mapping-timeout](#) on page 818
- [match-direction \(Services CoS\)](#) on page 818
- [match-direction \(Services IDS\)](#) on page 819
- [match-direction](#) on page 819
- [match-direction \(Services NAT\)](#) on page 820
- [match-direction \(Services Stateful Firewall\)](#) on page 820
- [max-drop-flows](#) on page 821
- [max-flows](#) on page 822
- [max-sessions-per-subscriber](#) on page 823
- [maximum](#) on page 823
- [maximum-contexts](#) on page 824
- [maximum-send-window](#) on page 824
- [member-failure-options \(Aggregated Multiservices\)](#) on page 825
- [member-interface \(Aggregated Multiservices\)](#) on page 827

- [message-rate-limit](#) on page 828
- [mlfr-uni-nni-bundles-inline](#) on page 829
- [mode](#) on page 830
- [mss](#) on page 830
- [multi-link-layer-2-inline](#) on page 831
- [multilink-class](#) on page 831
- [multilink-max-classes](#) on page 832
- [nat-options](#) on page 832
- [nat-rules](#) on page 833
- [next-hop-service](#) on page 834
- [no-anti-replay](#) on page 835
- [no-anti-replay \(Services Service Set\)](#) on page 835
- [no-fragmentation](#) on page 836
- [no-ipsec-tunnel-in-traceroute](#) on page 836
- [no-per-unit-scheduler](#) on page 837
- [no-termination-request](#) on page 837
- [no-translation](#) on page 838
- [output](#) on page 838
- [overload-pool](#) on page 839
- [overload-prefix](#) on page 839
- [passive-mode-tunneling](#) on page 840
- [pba-interim-logging-interval](#) on page 841
- [per-unit-scheduler](#) on page 842
- [perfect-forward-secrecy](#) on page 843
- [pgcp](#) on page 844
- [pgcp-rules](#) on page 844
- [policy \(Services IKE\)](#) on page 845
- [policy \(IPsec\)](#) on page 846
- [pool](#) on page 847
- [pool \(Service Interface\)](#) on page 848
- [port \(Services NAT\)](#) on page 849
- [port \(Services Voice\)](#) on page 851
- [port \(System Log Messsages\)](#) on page 851
- [port-forwarding](#) on page 852
- [port-forwarding-mappings](#) on page 852
- [ports-per-session](#) on page 853
- [post-service-filter](#) on page 853

- [ppp-access-profile](#) on page 854
- [pre-shared-key](#) (Services IKE) on page 854
- [preserve-interface](#) on page 855
- [primary](#) (Adaptive Services Interfaces) on page 855
- [primary](#) (Link Services IQ PIC Interfaces) on page 856
- [proposal](#) (Services IKE) on page 856
- [proposal](#) (Services IPsec VPN) on page 857
- [proposals](#) on page 857
- [protocol](#) (Applications) on page 858
- [protocol](#) (IPSec) on page 859
- [ptsp-rules](#) on page 859
- [queues](#) on page 860
- [reassembly-timeout](#) on page 860
- [receive-window](#) on page 861
- [redistribute-all-traffic](#) (Aggregated Multiservices) on page 861
- [redundancy-options](#) (Adaptive Services Interfaces) on page 862
- [redundancy-options](#) (Link Services IQ PIC Interfaces) on page 862
- [redundancy-options](#) (MS-MIC, MS-MPC) on page 863
- [\(reflexive | reverse\)](#) on page 864
- [rejoin-timeout](#) (Aggregated Multiservices) on page 865
- [remote-gateway](#) on page 865
- [remote-id](#) on page 866
- [remotely-controlled](#) on page 866
- [request-url](#) on page 867
- [replicate-services](#) (MS-MIC, MS-MPC) on page 868
- [respond-bad-spi](#) (Services IKE Policy) on page 869
- [retransmit-interval](#) (Services) on page 869
- [rpc-program-number](#) on page 870
- [routing-engine-services](#) on page 870
- [rtp](#) on page 871
- [rule](#) (Services CoS) on page 872
- [rule](#) (Services IDS) on page 873
- [rule](#) on page 875
- [rule](#) (Services NAT) on page 877
- [rule](#) (Services Stateful Firewall) on page 878
- [rule](#) (Softwire) on page 879
- [rule-set](#) (Services CoS) on page 879

- [rule-set \(Services IDS\) on page 880](#)
- [rule-set on page 880](#)
- [rule-set \(Services NAT\) on page 881](#)
- [rule-set \(Services Stateful Firewall\) on page 881](#)
- [rule-set \(Softwire\) on page 882](#)
- [secondary \(Adaptive Services Interfaces\) on page 882](#)
- [secondary \(Link Services IQ PIC Interfaces\) on page 883](#)
- [secure-nat-mapping on page 883](#)
- [secured-port-block-allocation on page 884](#)
- [server \(pcp\) on page 886](#)
- [service on page 887](#)
- [service-domain on page 888](#)
- [service-filter \(Interfaces\) on page 888](#)
- [service-interface \(Adaptive Services Interfaces\) on page 889](#)
- [service-interface \(L2TP Processing\) on page 889](#)
- [service-interface-pools on page 890](#)
- [service-set \(Interfaces\) on page 890](#)
- [service-set \(Services\) on page 891](#)
- [service-set-options on page 893](#)
- [services \(NAT\) on page 893](#)
- [session-limit on page 894](#)
- [set-dont-fragment-bit \(Services Set\) on page 895](#)
- [set-dont-fragment-bit \(Services IPsec VPN\) on page 895](#)
- [sip-call-hold-timeout on page 896](#)
- [sip on page 897](#)
- [snmp-command on page 897](#)
- [snmp-trap-thresholds on page 898](#)
- [softwire-concentrator on page 899](#)
- [softwire-options on page 900](#)
- [softwire-rules on page 900](#)
- [source-address \(Service Sets\) on page 901](#)
- [source-address \(Services CoS\) on page 901](#)
- [source-address \(Services IDS\) on page 902](#)
- [source-address on page 902](#)
- [source-address \(Services NAT\) on page 903](#)
- [source-address \(Services Stateful Firewall\) on page 903](#)
- [source-address-range \(Services IDS\) on page 904](#)



- [source-address-range \(Services NAT\) on page 904](#)
- [source-address-range \(Services Stateful Firewall\) on page 905](#)
- [source-pool on page 905](#)
- [source-port on page 906](#)
- [source-prefix \(Services IDS\) on page 906](#)
- [source-prefix \(Services NAT\) on page 907](#)
- [source-prefix-ipv6 on page 907](#)
- [source-prefix-list \(Services CoS\) on page 908](#)
- [source-prefix-list \(Services IDS\) on page 908](#)
- [source-prefix-list \(Services NAT\) on page 909](#)
- [source-prefix-list \(Services Stateful Firewall\) on page 909](#)
- [spi on page 910](#)
- [stateful-firewall-rules on page 910](#)
- [stateful-nat64 on page 911](#)
- [syslog \(Services CoS\) on page 911](#)
- [syslog \(Services IDS\) on page 912](#)
- [syslog on page 912](#)
- [syslog \(Services L2TP\) on page 913](#)
- [syslog \(Services NAT\) on page 913](#)
- [syslog \(Services Service Set\) on page 914](#)
- [syslog \(Services Stateful Firewall\) on page 915](#)
- [syn-cookie on page 916](#)
- [tcp-mss on page 917](#)
- [term \(Services CoS\) on page 918](#)
- [term \(Services IDS\) on page 919](#)
- [term on page 921](#)
- [term \(Services HCM\) on page 922](#)
- [term \(Services NAT\) on page 923](#)
- [term \(Services Stateful Firewall\) on page 924](#)
- [then \(Services CoS\) on page 925](#)
- [then \(Services HCM\) on page 925](#)
- [then \(Services IDS\) on page 926](#)
- [then on page 927](#)
- [then \(Services NAT\) on page 928](#)
- [then \(Services Stateful Firewall\) on page 929](#)
- [threshold \(Services IPsec\) on page 930](#)
- [threshold \(Services Logging and SYN-Cookie Defenses\) on page 930](#)

- [traceoptions \(Security PKI\) on page 931](#)
- [traceoptions \(Services IPsec VPN\) on page 933](#)
- [traceoptions \(Services L2TP\) on page 935](#)
- [traceoptions \(Services Logging\) on page 939](#)
- [translated on page 941](#)
- [transport on page 941](#)
- [trigger-link-failure on page 942](#)
- [translated-port on page 942](#)
- [translation-type on page 943](#)
- [trusted-ca on page 945](#)
- [ttl-threshold on page 945](#)
- [tunnel-group on page 946](#)
- [tunnel-mtu \(Services IPsec VPN\) on page 947](#)
- [tunnel-mtu \(Services Service Set\) on page 948](#)
- [tunnel-timeout on page 949](#)
- [url on page 949](#)
- [url-list on page 950](#)
- [url-rule on page 950](#)
- [url-rule-set on page 951](#)
- [unit \(Aggregated Multiservices\) on page 951](#)
- [unit \(Interfaces\) on page 952](#)
- [unit \(Voice Services\) on page 953](#)
- [uuid on page 954](#)
- [v6rd on page 955](#)
- [version \(IKE\) on page 956](#)
- [video on page 956](#)
- [video \(Application Profile\) on page 957](#)
- [voice on page 957](#)
- [voice \(Application Profile\) on page 958](#)
- [warm-standby on page 958](#)

---

## [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*
  - *Application Identification for Nested Applications*
  - *Configuring Global APPID Properties*

## 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 415](#)
- [Configuring IKE Proposals on page 435](#)
- [Configuring IKE Policies on page 439](#)
- [Configuring IPsec Proposals on page 445](#)
- [Configuring IPsec Policies on page 450](#)
- [Configuring IPsec Rules on page 452](#)
- [Configuring IPsec Rule Sets on page 459](#)
- [Configuring Dynamic Endpoints for IPsec Tunnels on page 495](#)
- [Tracing Junos VPN Site Secure Operations on page 466](#)

- [Configuring Junos VPN Site Secure or IPSec VPN on page 549](#)

## adaptive-services-pics

<b>Syntax</b>	<pre>adaptive-services-pics {   traceoptions {     file filename &lt;files number&gt; &lt;match regular-expression&gt; &lt;size size&gt; &lt;world-readable         no-world-readable&gt;;     flag flag;     no-remote-trace;   } }</pre>
<b>Hierarchy Level</b>	[edit services]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. The <b>file</b> option was added in Release 8.0.
<b>Description</b>	Define global services properties.
<b>Options</b>	The remaining statement is explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Tracing Services PIC Operations on page 27</a></li> </ul>

## address (Interfaces)

---

<b>Syntax</b>	<code>address address {     ... }</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>family</b> <i>family</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>family</b> <i>family</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the interface address.
<b>Options</b>	<b>address</b> —Address of the interface.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces.</li><li>• <a href="#">Configuring the Logical Interface Address for the MLPPP Bundle on page 666</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>

## address (Services NAT Pool)

---

<b>Syntax</b>	<code>address ip-prefix&lt;/prefix-length&gt;;</code>
<b>Hierarchy Level</b>	[edit <b>services</b> nat <b>pool</b> <i>nat-pool-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>prefix</b> option enhanced to support IPv6 addresses in Junos OS Release 8.5.
<b>Description</b>	Specify the NAT pool prefix value.
<b>Options</b>	<b>prefix</b> —Specify an IPv4 or IPv6 prefix value.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Source and Destination Addresses Network Address Translation Overview on page 52</a></li></ul>



## address-allocation

<b>Syntax</b>	address-allocation round-robin;
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">pool</a> <i>pool-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2.
<b>Description</b>	<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 <b>address-allocation round-robin</b> statement, round-robin allocation is enabled by default on MS-MICs and MS-MPCs.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Pools of Addresses and Ports for Network Address Translation Overview on page 53</a></li> </ul>

## address-range

<b>Syntax</b>	address-range low <i>minimum-value</i> high <i>maximum-value</i> ;
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">pool</a> <i>nat-pool-name</i> ]
<b>Release Information</b>	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.
<b>Description</b>	Specify the NAT pool address range.
<b>Options</b>	<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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Source and Destination Addresses Network Address Translation Overview on page 52</a></li> </ul>

## aggregation

---

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit services ids <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the type of data to be aggregated.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IDS Rules on page 384</a></li></ul>

## 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

**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 Stateful Firewall Rules on page 362](#)

## allow-multicast

---

<b>Syntax</b>	allow-multicast;
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.0.
<b>Description</b>	Allow multicast traffic to be sent to the Adaptive Services or Multiservices PIC.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Enabling Services PICs to Accept Multicast Traffic on page 17</a></li></ul>

## allow-overlapping-nat-pools

---


<b>Syntax</b>	allow-overlapping-nat-pools;
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat]
<b>Release Information</b>	Statement introduced with Junos OS Release 12.1.
<b>Description</b>	Specify that NAT source or destination pools can be shared between multiple service sets.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Sets for Network Address Translation on page 61</a></li></ul>

## anti-replay-window-size (Services IPsec VPN)

---

<b>Syntax</b>	anti-replay-window-size <i>bits</i> ;
<b>Hierarchy Level</b>	[edit services ipsec-vpn <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0.
<b>Description</b>	Specify the size of the IPsec antireplay window.
<b>Options</b>	<b>bits</b> —Size of the antireplay window, in bits. <b>Default:</b> 64 bits (AS PICs), 128 bits (Multiservices PICs and DPCs) <b>Range:</b> 64 through 4096 bits
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.

## anti-replay-window-size (Services Service Set)

<b>Syntax</b>	<code>anti-replay-window-size <i>bits</i>;</code>
<b>Hierarchy Level</b>	<code>[edit services service-set <i>service-set-name</i> ipsec-vpn-options]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0.
<b>Description</b>	<p>Specify the size of the IPsec antireplay window. This statement is useful for dynamic endpoint tunnels for which you cannot configure the <b>anti-replay-window-size</b> statement at the <code>[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then]</code> 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 <b>anti-replay-window-size</b> statement at the <code>[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then]</code> hierarchy level. If antireplay check has to be disabled for a particular tunnel in this service set, set the <b>no-anti-replay</b> statement at the <code>[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then]</code> hierarchy level.</p>
	<p> <b>NOTE:</b> The <b>anti-replay-window-size</b> and <b>no-anti-replay</b> settings at the <code>[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then]</code> hierarchy level override the settings specified at the <code>[edit services service-set <i>service-set-name</i> ipsec-vpn-options]</code> hierarchy level.</p>
<b>Options</b>	<p><b>bits</b>—Size of the antireplay window, in bits.</p> <p><b>Default:</b> 64 bits (AS PICs), 128 bits (Multiservices PICs and DPCs)</p> <p><b>Range:</b> 64 through 4096 bits</p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPsec Service Sets on page 460</a></li> <li>• <a href="#">Configuring IPsec Rules on page 452</a></li> </ul>

## app-mapping-timeout

---

<b>Syntax</b>	<code>app-mapping-timeout <i>app-mapping-timeout</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">pool</a> <i>nat-pool-name</i> ]
<b>Release Information</b>	<code>mapping-timeout</code> statement introduced in JUNOS Release 12.3.
<b>Description</b>	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 <a href="#">mapping-timeout</a> , the timeout value configured for <a href="#">mapping-timeout</a> is used. If neither option is specified, the default value of 300 seconds is used.
<b>Options</b>	<b><code>app-mapping-timeout</code></b> —Lifetime of AP-P mappings in seconds. <b>Default:</b> 300 <b>Range:</b> 120 through 864,000
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Source and Destination Addresses Network Address Translation Overview on page 52</a></li></ul>

## application

---

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit <a href="#">applications</a> ], [edit <a href="#">applications application-set</a> <i>application-set-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure properties of an application and whether to include it in an application set.
<b>Options</b>	<p><b><i>application-name</i></b>—Identifier of the application.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">ALG Descriptions on page 299</a></li> <li>• <a href="#">Configuring Application Sets on page 325</a></li> <li>• <a href="#">Configuring Application Protocol Properties on page 325</a></li> <li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li> <li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li> </ul>

## application-protocol

---

<b>Syntax</b>	<code>application-protocol <i>protocol-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">applications</a> <a href="#">application</a> <i>application-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>login</b> options introduced in Junos OS Release 7.4. <b>ip</b> option introduced in Junos OS Release 8.2.
<b>Description</b>	Identify the application protocol name. Application protocols are also called application layer gateways (ALGs).
<b>Options</b>	<p><b><i>protocol-name</i></b>—Name of the protocol. The following protocols are supported:</p> <ul style="list-style-type: none"><li><b>bootp</b>—Bootstrap protocol</li><li><b>dce-rpc</b>—DCE RPC</li><li><b>dce-rpc-portmap</b>—DCE RPC portmap</li><li><b>dns</b>—Domain Name Service</li><li><b>exec</b>—Remote Execution Protocol</li><li><b>ftp</b>—File Transfer Protocol</li><li><b>h323</b>—H.323</li><li><b>icmp</b>—ICMP</li><li><b>iiop</b>—Internet Inter-ORB Protocol</li><li><b>ip</b>—IP</li><li><b>login</b>—Login</li><li><b>netbios</b>—NetBIOS</li><li><b>netshow</b>—NetShow</li><li><b>pptp</b>—Point-to-Point Tunneling Protocol</li><li><b>realaudio</b>—RealAudio</li><li><b>rpc</b>—RPC</li><li><b>rpc-portmap</b>—RPC portmap</li><li><b>rtsp</b>—Real Time Streaming Protocol</li><li><b>shell</b>—Shell</li><li><b>sip</b>—Session Initiation Protocol</li><li><b>snmp</b>—SNMP</li><li><b>sqlnet</b>—SQLNet</li><li><b>talk</b>—Talk Program</li></ul>



**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 299](#)
- [Configuring Application Sets on page 325](#)
- [Configuring Application Protocol Properties on page 325](#)
- [Examples: Configuring Application Protocols on page 343](#)
- [Verifying the Output of ALG Sessions on page 344](#)

## 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 <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> ], [edit services cos <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> ( <b>reflexive</b>   <b>reverse</b> )]
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	<b><i>profile-name</i></b> —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"><li><a href="#">Configuring Application Profiles for Use as CoS Rule Actions on page 559</a></li></ul>

## application-set

---

<b>Syntax</b>	<code>application-set <i>application-set-name</i> {     <a href="#">application</a> <i>application-name</i>; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">applications</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure one or more applications to include in an application set.
<b>Options</b>	<i>application-set-name</i> —Identifier of an application set.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">ALG Descriptions on page 299</a></li> <li>• <a href="#">Configuring Application Sets on page 325</a></li> <li>• <a href="#">Configuring Application Protocol Properties on page 325</a></li> <li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li> <li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li> </ul>

## application-sets (Services CoS)

---

<b>Syntax</b>	<code>applications-sets <i>set-name</i>;</code>
<b>Hierarchy Level</b>	[edit services cos <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1.
<b>Description</b>	Define one or more target application sets.
<b>Options</b>	<i>set-name</i> —Name of the target application set.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Match Conditions In CoS Rules on page 557</a></li> </ul>

## application-sets (Services IDS)

---

<b>Syntax</b>	<code>application-sets set-name;</code>
<b>Hierarchy Level</b>	[edit services ids <a href="#">rule rule-name</a> <a href="#">term term-name</a> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define one or more target application sets.
<b>Options</b>	<i>set-name</i> —Name of the target application set.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Match Conditions in IDS Rules on page 386</a></li></ul>

## application-sets (Services NAT)

---

<b>Syntax</b>	<code>applications-sets set-name;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule rule-name</a> <a href="#">term term-name</a> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define one or more target application sets.
<b>Options</b>	<i>set-name</i> —Name of the target application set.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li></ul>

## application-sets (Services Stateful Firewall)

---

<b>Syntax</b>	<code>applications-sets <i>set-name</i>;</code>
<b>Hierarchy Level</b>	[edit services stateful-firewall <a href="#">rule rule-name</a> <a href="#">term term-name</a> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define one or more target application sets.
<b>Options</b>	<i>set-name</i> —Name of the target application set.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Match Conditions in Stateful Firewall Rules on page 360</a></li> </ul>

## applications (Services ALGs)

---

<b>Syntax</b>	<code>applications { ... }</code>
<b>Hierarchy Level</b>	[edit]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define the applications used in services.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">ALG Descriptions on page 299</a></li> <li>• <a href="#">Configuring Application Sets on page 325</a></li> <li>• <a href="#">Configuring Application Protocol Properties on page 325</a></li> <li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li> <li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li> </ul>

## applications (Services CoS)

---

<b>Syntax</b>	<code>applications [ <i>application-name</i> ];</code>
<b>Hierarchy Level</b>	[edit services cos <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>from</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1.
<b>Description</b>	Define one or more applications to which the CoS services apply.
<b>Options</b>	<i>application-name</i> —Name of the target application.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Match Conditions in a CoS Rule</a></li><li>• <a href="#">Configuring Match Conditions In CoS Rules on page 557</a></li></ul>

## applications (Services IDS)

---

<b>Syntax</b>	<code>applications [ <i>application-name</i> ];</code>
<b>Hierarchy Level</b>	[edit services ids <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>from</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define one or more applications to which IDS applies.
<b>Options</b>	<i>application-name</i> —Name of the target application.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Match Conditions in IDS Rules on page 386</a></li></ul>

## applications (Services NAT)

---

<b>Syntax</b>	<code>applications [ <i>application-name</i> ];</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define one or more application protocols to which the NAT services apply.
<b>Options</b>	<i>application-name</i> —Name of the target application.
<b>Required Privilege</b>	interface—To view this statement in the configuration.
<b>Level</b>	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> </ul>

## applications (Services Stateful Firewall)

---

<b>Syntax</b>	<code>applications [ <i>application-name</i> ];</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> stateful-firewall <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define one or more applications to which the stateful firewall services apply.
<b>Options</b>	<i>application-name</i> —Name of the target application.
<b>Required Privilege</b>	interface—To view this statement in the configuration.
<b>Level</b>	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Match Conditions in Stateful Firewall Rules on page 360</a></li> </ul>

## authentication

---

<b>Syntax</b>	<pre>authentication {   algorithm (hmac-md5-96   hmac-sha1-96);   key (ascii-text key   hexadecimal key); }</pre>
<b>Hierarchy Level</b>	[edit services (IPsec VPN) ipsec-vpn <a href="#">rule (Services IPsec VPN) rule-name</a> <a href="#">term (Services IPsec VPN) term-name</a> <a href="#">then (Services IPsec VPN) manual direction direction</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure IPsec authentication parameters for a manual security association (SA).
<b>Options</b>	<p><b>algorithm</b>—Hash algorithm that authenticates packet data. The algorithm can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>hmac-md5-96</b>—Produces a 128-bit digest.</li><li>• <b>hmac-sha1-96</b>—Produces a 160-bit digest.</li></ul> <p><b>key</b>—Type of authentication key. The key can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>ascii-text key</b>—ASCII text key. For <b>hmac-md5-96</b>, the key is 16 ASCII characters; for <b>hmac-sha1-96</b>, the key is 20 ASCII characters.</li><li>• <b>hexadecimal key</b>—Hexadecimal key. For <b>hmac-md5-96</b>, the key is 32 hexadecimal characters; for <b>hmac-sha1-96</b>, the key is 40 hexadecimal characters.</li></ul>
<b>Usage Guidelines</b>	See <i>Configuring Security Associations</i> .
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>



---

## authentication-algorithm (Services IKE)

---

<b>Syntax</b>	authentication-algorithm (md5   sha1   sha-256);
<b>Hierarchy Level</b>	[edit services ipsec-vpn <a href="#">ike proposal</a> <i>proposal-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. sha-256 option added in Junos OS Release 7.6.
<b>Description</b>	Configure the Internet Key Exchange (IKE) hash algorithm that authenticates packet data.
<b>Options</b>	<b>md5</b> —Produces a 128-bit digest. <b>sha1</b> —Produces a 160-bit digest. <b>sha-256</b> —Produces a 256-bit digest. <b>sha-384</b> —Produces a 384-bit digest.
<b>Required Privilege Level</b>	<b>admin</b> —To view this statement in the configuration. <b>admin-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IKE Proposals on page 435</a></li></ul>

## authentication-algorithm (Services IPsec)

<b>Syntax</b>	authentication-algorithm (hmac-md5-96   hmac-sha-256-128   hmac-sha1-96);
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>ipsec proposal</b> <i>ipsec-proposal-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.

---

<b>Options</b>	<p><b>hmac-md5-96</b>—Produces a 128-bit digest.</p> <p><b>hmac-sha-256-128</b>—Produces a 256-bit digest.</p> <p><b>hmac-sha1-96</b>—Produces a 160-bit digest.</p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPsec Proposals on page 445</a></li> </ul>

## authentication-method

---

<b>Syntax</b>	authentication-method (dsa-signatures   pre-shared-keys   rsa-signatures);
<b>Hierarchy Level</b>	[edit services (IPsec VPN) ipsec-vpn <b>ike proposal</b> <i>proposal-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure an IKE authentication method.
<b>Options</b>	<p><b>dsa-signatures</b>—Digital signature algorithm (DSA).</p> <p><b>rsa-signatures</b>—Public key algorithm (supports encryption and digital signatures).</p> <p><b>pre-shared-keys</b>—A key derived from an out-of-band mechanism; the key authenticates the exchange.</p>
<b>Usage Guidelines</b>	See <i>Configuring IKE Proposals</i> .
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>

## auxiliary-spi

---

<b>Syntax</b>	<code>auxiliary-spi spi-value;</code>
<b>Hierarchy Level</b>	[edit services (IPsec VPN) ipsec-vpn <b>rule</b> (Services IPsec VPN) <b>rule-name</b> <b>term</b> (Services IPsec VPN) <b>term-name</b> <b>then</b> (Services IPsec VPN) <b>manual direction</b> <b>direction</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure an auxiliary Security Parameter Index (SPI) for a manual SA. Use the auxiliary SPI when you configure the <b>protocol</b> statement to use the <b>bundle</b> option.
<b>Options</b>	<b>spi-value</b> —An arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet). <b>Range:</b> 256 through 16,639
<b>Usage Guidelines</b>	See <i>Configuring Security Associations</i> . For information about SPI, see <i>Configuring Security Associations</i> and <b>spi</b> .
<b>Required Privilege Level</b>	<b>admin</b> —To view this statement in the configuration. <b>admin-control</b> —To add this statement to the configuration.

## backup-remote-gateway

---

<b>Syntax</b>	<code>backup-remote-gateway address;</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>rule</b> <b>rule-name</b> <b>term</b> <b>term-name</b> <b>then</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<b>address</b> —Backup remote IPv4 or IPv6 address.
<b>Required Privilege Level</b>	<b>admin</b> —To view this statement in the configuration. <b>admin-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IPsec Rules on page 452</a></li></ul>

## bundle

<b>Syntax</b>	<code>bundle (lsq-fpc/pic/port   ... );</code>
<b>Hierarchy Level</b>	[edit interfaces lsq-fpc/pic/port <b>unit</b> logical-unit-number <b>family</b> mlppp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Associate the voice services interface with the logical interface it is joining.
<b>Options</b>	<b>lsq-fpc/pic/port</b> —Name of the voice services interface you are linking.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Voice Services Bundles with MLPPP Encapsulation on page 670</a></li> </ul>

## by-destination

<b>Syntax</b>	<pre>by-destination {   hold-time <i>seconds</i>;   maximum <i>number</i>;   packets <i>number</i>;   rate <i>number</i>; }</pre>
<b>Hierarchy Level</b>	[edit services ids <b>rule</b> rule-name <b>term</b> term-name <b>then</b> session-limit]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Apply limit to sessions based on numbers generated from the configured destination (IP or subnet) or application.
<b>Options</b>	<p><b>hold-time <i>seconds</i></b>—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 <b>maximum</b>, <b>packets</b>, or <b>rate</b> statements.</p> <p><b>maximum <i>number</i></b>—Maximum number of open sessions per application or IP address.</p> <p><b>packets <i>number</i></b>—Maximum peak packets per second per application or IP address.</p> <p><b>rate <i>number</i></b>—Maximum number of sessions per second per application or IP address.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Actions in IDS Rules on page 387</a></li> </ul>

## 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 <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> <i>session-limit</i> ]
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><b>hold-time <i>seconds</i></b>—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 <b>maximum</b>, <b>packets</b>, or <b>rate</b> statements.</p> <p><b>maximum <i>number</i></b>—Maximum number of open sessions per application or IP address.</p> <p><b>packets <i>number</i></b>—Maximum peak packets per second per application or IP address.</p> <p><b>rate <i>number</i></b>—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"><li>• <a href="#">Configuring Actions in IDS Rules on page 387</a></li></ul>

## by-source

<b>Syntax</b>	by-source { hold-time <i>seconds</i> ; maximum <i>number</i> ; packets <i>number</i> ; rate <i>number</i> ; }
<b>Hierarchy Level</b>	[edit services ids <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> <i>session-limit</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Apply limit to sessions based on numbers generated from the configured source (IP or subnet) or application.
<b>Options</b>	<p><b>hold-time <i>seconds</i></b>—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 <b>maximum</b>, <b>packets</b>, or <b>rate</b> statements.</p> <p><b>maximum <i>number</i></b>—Maximum number of open sessions per application or IP address.</p> <p><b>packets <i>number</i></b>—Maximum peak packets per second per application or IP address.</p> <p><b>rate <i>number</i></b>—Maximum number of sessions per second per application or IP address.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Actions in IDS Rules on page 387</a></li> </ul>

## bypass-traffic-on-exceeding-flow-limits

<b>Syntax</b>	bypass-traffic-on-exceeding-flow-limits;
<b>Hierarchy Level</b>	[edit services service-set <i>service-set-name</i> service-set-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	Enable packets to bypass without creating a new session when the flow in the service set exceeds the limit that is set by the <b>max-flows</b> statement at the [edit services service-set <i>service-set-name</i> ] hierarchy level.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Service Sets to be Applied to Services Interfaces on page 9</a></li> </ul>

## bypass-traffic-on-pic-failure

---

<b>Syntax</b>	bypass-traffic-on-pic-failure;
<b>Hierarchy Level</b>	[edit services service-set <i>service-set-name</i> service-set-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	<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 <b>bypass-traffic-on-pic-failure</b> 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>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Sets to be Applied to Services Interfaces on page 9</a></li></ul>

## cgn-pic

---

<b>Syntax</b>	cgn-pic;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> services-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2.
<b>Description</b>	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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 41</a></li></ul>



## cisco-interopability

---

<b>Syntax</b>	<code>cisco-interopability send-lip-remove-link-for-link-reject;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4.
<b>Description</b>	FRF.16 interoperability settings.
<b>Options</b>	<b>send-lip-remove-link-for-link-reject</b> —Send Link Integrity Protocol remove link when an add-link rejection message is received.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring SONET APS Interoperability with Cisco Systems FRF.16 on page 591</a></li></ul>

## class

---

<b>Syntax</b>	<pre>class {     alg-logs;     ids-logs;     nat-logs;     packet-logs;     pcp-logs;     session-logs &lt;open   close&gt;;     stateful-firewall-logs ; }</pre>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> <a href="#">syslog host</a> <i>hostname</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2.
<b>Description</b>	Set the class of applications to be logged to the system log.
<b>Options</b>	<p><i>class-name</i>—Enter one of the following values:</p> <ul style="list-style-type: none"><li>• <b>alg-logs</b>—Log application-level gateway events.</li><li>• <b>ids-logs</b>—Log intrusion detection system events.</li><li>• <b>nat-logs</b>—Log Network Address Translation events.</li><li>• <b>packet-logs</b>—Log general packet-related events.</li><li>• <b>session-logs</b>—Log session open and close events.</li><li>• <b>session-logs open</b>—Log session open events only.</li><li>• <b>session-logs close</b>—Log session close events.</li></ul>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• See <a href="#">Configuring System Logging for Service Sets on page 26</a>.</li></ul>

## clear-dont-fragment-bit (Interfaces GRE Tunnels)

<b>Syntax</b>	clear-dont-fragment-bit;
<b>Hierarchy Level</b>	[edit interfaces gr-fpc/pic/port <b>unit</b> logical-unit-number], [edit logical-systems logical-system-name interfaces gr-fpc/pic/port <b>unit</b> logical-unit-number]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<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 <b>clear-dont-fragment-bit</b> 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>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Enabling Fragmentation on GRE Tunnels</i></li> </ul>

## clear-dont-fragment-bit

<b>Syntax</b>	clear-dont-fragment-bit;
<b>Hierarchy Level</b>	[edit services (IPsec VPN) ipsec-vpn <b>rule</b> (Services IPsec VPN) rule-name <b>term</b> (Services IPsec VPN) term-name <b>then</b> (Services IPsec VPN)]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Usage Guidelines</b>	See <i>Configuring IPsec Rules</i> .
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## clear-dont-fragment-bit (Services NAT Options)

---

<b>Syntax</b>	clear-dont-fragment-bit;
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> <a href="#">nat-options</a> <a href="#">stateful-nat64</a> ]
<b>Release Information</b>	Statement introduced with Junos OS Release 12.1.
<b>Description</b>	Clear the DF (don't fragment) bit in a translated IPv4 packet if its packet size is less than 1280 bytes. If the packet is greater than or equal to 1280 bytes, the DF bit is not cleared.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Rules on page 14</a></li></ul>

## clear-dont-fragment-bit (Services Service Set)

---

<b>Syntax</b>	clear-dont-fragment-bit;
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> <a href="#">ipsec-vpn-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0.
<b>Description</b>	<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 <b>clear-dont-fragment-bit</b> statement at the [edit services <a href="#">ipsec-vpn</a> 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 <b>clear-dont-fragment-bit</b> statement at the [edit services <a href="#">ipsec-vpn</a> rule <i>rule-name</i> term <i>term-name</i> then] hierarchy level.</p> <p>By default, this statement is disabled (the DF bit value is not cleared on the inner header and outer header by default).</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IPsec Service Sets on page 460</a></li><li>• <a href="#">Configuring IPsec Rules on page 452</a></li></ul>

## clear-ike-sas-on-pic-restart

---

<b>Syntax</b>	clear-ike-sas-on-pic-restart;
<b>Hierarchy Level</b>	[edit services ipsec-vpn]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5.
<b>Description</b>	Clear IKE security associations (SAs) when the corresponding PIC restarts or is taken offline.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Security Associations on page 415</a></li></ul>

## clear-ipsec-sas-on-pic-restart

---

<b>Syntax</b>	clear-ipsec-sas-on-pic-restart;
<b>Hierarchy Level</b>	[edit services ipsec-vpn]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2.
<b>Description</b>	Clear IPsec security associations (SAs) when the corresponding PIC restarts or is taken offline.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Security Associations on page 415</a></li></ul>

## compression

---

<b>Syntax</b>	<pre>compression {   rtp {     f-max-period <i>number</i>;     maximum-contexts <i>number</i> &lt;force&gt;;     port {       minimum <i>port-number</i>;       maximum <i>port-number</i>;     }     queues [ <i>queue-numbers</i> ];   } }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the compression properties for voice services traffic.  The remaining statements are described separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Compression of Voice Traffic on page 667</a></li></ul>

## compression-device (Interfaces)

---

<b>Syntax</b>	<pre>compression-device <i>interface-name</i>;</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Specify the compression interface for voice services traffic.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Compression Interface with PPP Encapsulation on page 670</a></li></ul>

## copy-dont-fragment-bit (Services IPsec VPN)

<b>Syntax</b>	<code>copy-dont-fragment-bit;</code>
<b>Hierarchy Level</b>	<code>[edit services ipsec-vpn <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	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 <b>copy-dont-fragment-bit</b> statement at the <code>[edit services service-set <i>service-set-name</i> ipsec-vpn-options]</code> 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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## copy-dont-fragment-bit (Services Set)

<b>Syntax</b>	<code>copy-dont-fragment-bit;</code>
<b>Hierarchy Level</b>	<code>[edit services service-set <i>service-set-name</i> ipsec-vpn-options]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	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 <b>copy-dont-fragment-bit</b> statement at the <code>[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> <b>then</b>]</code> 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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPsec Service Sets on page 460</a></li> <li>• <a href="#">Configuring IPsec Rules on page 452</a></li> </ul>

## data (FTP)

---

<b>Syntax</b>	<pre>data {     dscp (<i>alias</i>   <i>bits</i>);     forwarding-class <i>class-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit services cos <a href="#">application-profile</a> <i>profile-name</i> ftp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3.
<b>Description</b>	Set the appropriate <b>dscp</b> and <b>forwarding-class</b> value for FTP data.
<b>Default</b>	By default, the system will not alter the DSCP or forwarding class for FTP data traffic.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Application Profiles</i></li><li>• <a href="#">video (Application Profile) on page 957</a></li><li>• <a href="#">voice (Application Profile) on page 958</a></li></ul>

## dead-peer-detection (Services IPsec VPN)

---

<b>Syntax</b>	<pre>dead-peer-detection {     <a href="#">interval</a> <i>seconds</i>;     <a href="#">threshold</a> <i>number</i>; }</pre>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <i>rule-name</i> term <i>term-name</i> then]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	Sets dead peer detection options when dead peer detection has been enabled with the <a href="#">initiate-dead-peer-detection</a> command. The <b>dead-peer-detection</b> options are used for IKEv1 security associations (SAs) but not for IKEv2 SAs.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	security—To view this statement in the configuration. security-control—To add this statement to the configuration.



## description (Services IPsec VPN)

<b>Syntax</b>	<code>description <i>description</i>;</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>ike</b> <b>policy</b> <i>policy-name</i> ], [edit services ipsec-vpn <b>ike</b> <b>proposal</b> <i>proposal-name</i> ], [edit services ipsec-vpn <b>ipsec</b> <b>policy</b> <i>policy-name</i> ], [edit services ipsec-vpn <b>ipsec</b> <b>proposal</b> <i>proposal-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the text description for an IKE or IPsec policy or proposal.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">description on page 743</a></li> <li>• <a href="#">Configuring IPsec Proposals on page 445</a></li> <li>• <a href="#">Configuring IPsec Policies on page 450</a></li> </ul>

## destination-address (Services CoS)

<b>Syntax</b>	<code>destination-address (<i>address</i>   any-unicast) &lt;except&gt;;</code>
<b>Hierarchy Level</b>	[edit services cos <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>from</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. <b>address</b> option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.
<b>Description</b>	Specify the destination address for rule matching.
<b>Options</b>	<b>address</b> —Destination IPv4 or IPv6 address or prefix value.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Match Conditions in a CoS Rule</a></li> <li>• <a href="#">Configuring Match Conditions In CoS Rules on page 557</a></li> </ul>

## destination-address (Services IDS)

---

<b>Syntax</b>	<code>destination-address (address   any-unicast) &lt;except&gt;;</code>
<b>Hierarchy Level</b>	[edit services ids <a href="#">rule rule-name term term-name from</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>address</b> option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.
<b>Description</b>	Specify the destination address for rule matching.
<b>Options</b>	<b>address</b> —Destination IPv4 or IPv6 address or prefix value. <b>any-unicast</b> —Any unicast packet. <b>except</b> —(Optional) Exempt the specified address, prefix, or unicast packets from rule matching.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Match Conditions in IDS Rules on page 386</a></li></ul>

## destination-address

---

<b>Syntax</b>	<code>destination-address address;</code>
<b>Hierarchy Level</b>	[edit services (IPsec VPN) ipsec-vpn <a href="#">rule (Services IPsec VPN) rule-name term (Services IPsec VPN) term-name from (Services IPsec VPN)</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the destination address for rule matching.
<b>Options</b>	<b>address</b> —Destination IP address.
<b>Usage Guidelines</b>	See <i>Configuring IPsec Rules</i> .
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## destination-address (Services NAT)

<b>Syntax</b>	<code>destination-address (address   any-unicast) &lt;except&gt;;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> rule-name <a href="#">term</a> term-name <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>any-unicast</b> and <b>except</b> options introduced in Junos OS Release 7.6. <b>address</b> option enhanced to support IPv6 and addresses in Junos OS Release 8.5.
<b>Description</b>	Specify the destination address for rule matching.
<b>Options</b>	<b>address</b> —Destination IPv4 or IPv6 address or prefix value.  <b>any-unicast</b> —Any unicast packet.  <b>except</b> —(Optional) Prevent the specified address, prefix, or unicast packets from being translated.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> </ul>

## destination-address (Services Stateful Firewall)

<b>Syntax</b>	<code>destination-address (address   any-unicast) &lt;except&gt;;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> stateful-firewall <a href="#">rule</a> rule-name <a href="#">term</a> term-name <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>any-unicast</b> and <b>except</b> options introduced in Junos OS Release 7.6. <b>address</b> option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.
<b>Description</b>	Specify the destination address for rule matching.
<b>Options</b>	<b>address</b> —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.  <b>any-unicast</b> —Match all unicast packets.  <b>except</b> —(Optional) Exclude the specified address, prefix, or unicast packets from rule matching.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Match Conditions in Stateful Firewall Rules on page 360</a></li> </ul>

## destination-address-range (Services IDS)

---

<b>Syntax</b>	<code>destination-address-range low <i>minimum-value</i> high <i>maximum-value</i> &lt;except&gt;;</code>
<b>Hierarchy Level</b>	[edit services ids <a href="#">rule rule-name</a> <a href="#">term term-name</a> <a href="#">from</a> ]
<b>Release Information</b>	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.
<b>Description</b>	Specify the destination address range for rule matching.
<b>Options</b>	<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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Match Conditions in IDS Rules on page 386</a></li></ul>

## destination-address-range (Services NAT)

<b>Syntax</b>	<code>destination-address-range low <i>minimum-value</i> high <i>maximum-value</i> &lt;except&gt;;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">from</a> ]
<b>Release Information</b>	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.
<b>Description</b>	Specify the destination address range for rule matching.  If the <a href="#">translation-type</a> statement in the <a href="#">then</a> statement of the nat rule is set to <b>stateful-nat-64</b> , the destination address range for rule matching must be within the range specified by the <a href="#">destination-prefix</a> statement in the <a href="#">then</a> statement.
<b>Options</b>	<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.  <b>except</b> —(Optional) Prevent the specified address range from being translated.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> </ul>

## destination-address-range (Services Stateful Firewall)

---

<b>Syntax</b>	<code>destination-address-range low <i>minimum-value</i> high <i>maximum-value</i> &lt;except&gt;;</code>
<b>Hierarchy Level</b>	[edit services stateful-firewall <a href="#">rule rule-name</a> <a href="#">term term-name</a> <a href="#">from</a> ]
<b>Release Information</b>	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.
<b>Description</b>	Specify the destination address range for rule matching.
<b>Options</b>	<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 range from rule matching.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Match Conditions in Stateful Firewall Rules on page 360</a></li></ul>

## destination-pool

---

<b>Syntax</b>	<code>destination-pool <i>nat-pool-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule rule-name</a> <a href="#">term term-name</a> <a href="#">then translated</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the destination address pool for translated traffic.
<b>Options</b>	<i>nat-pool-name</i> —Destination pool name.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li></ul>

## destination-port

<b>Syntax</b>	<code>destination-port <i>port-value</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">applications application application-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number.
<b>Options</b>	<p><b><i>port-value</i></b>—Identifier for the port or range of ports. For a complete list of supported application destination port requirements, see <a href="#">“Configuring Source and Destination Ports” on page 331</a>.</p> <p><b>Range:</b> 1 through 65,535</p>



**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
```

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">ALG Descriptions on page 299</a></li> <li>• <a href="#">Configuring Application Sets on page 325</a></li> <li>• <a href="#">Configuring Application Protocol Properties on page 325</a></li> <li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li> <li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li> <li>• <a href="#">Two-Way Active Measurement Protocol Overview</a></li> </ul>

## destination-port range

---

<b>Syntax</b>	<code>destination-port range <i>high</i>   <i>low</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services nat rule rule-name term term-name from</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	Specify the destination port range for rule matching.
<b>Options</b>	<i>high</i> —Upper limit of port range for matching. <i>low</i> —Lower limit of port range for matching.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Port Forwarding for Static Destination Address Translation on page 173</a></li></ul>

## destination-prefix (Services IDS)

---

<b>Syntax</b>	<code>destination-prefix <i>prefix-value</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services ids rule rule-name term term-name then aggregation</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the prefix value for destination IPv4 address aggregation.
<b>Options</b>	<i>prefix-value</i> —Integer value. <b>Range:</b> 1 through 32
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Actions in IDS Rules on page 387</a></li></ul>



## destination-prefix (Services NAT)

---

<b>Syntax</b>	<code>destination-prefix <i>destination-prefix</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then translated</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. <i>destination-prefix</i> option enhanced to support IPv6 addresses in Junos OS Release 8.5.
<b>Description</b>	Specify the destination prefix for translated traffic.
<b>Options</b>	<i>destination-prefix</i> —IPv4 or IPv6 destination prefix value.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> </ul>

## destination-prefix-ipv6

---

<b>Syntax</b>	<code>destination-prefix-ipv6 <i>prefix</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> ids <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then aggregation</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5.
<b>Description</b>	Specify the prefix value for destination IPv6 address aggregation.
<b>Options</b>	<i>prefix-value</i> —Integer value. <b>Range:</b> 1 through 128
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Actions in IDS Rules on page 387</a></li> </ul>

## destination-prefix-list (Services CoS)

---

<b>Syntax</b>	<code>destination-prefix-list <i>list-name</i> &lt;except&gt;;</code>
<b>Hierarchy Level</b>	<code>[edit services cos <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>from</b>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2.
<b>Description</b>	Specify the destination prefix list for rule matching. You configure the prefix list by including the <b>prefix-list</b> statement at the <b>[edit policy-options]</b> hierarchy level.
<b>Options</b>	<b><i>list-name</i></b> —Destination prefix list.  <b>except</b> —(Optional) Exclude the specified prefix list from rule matching.
<b>Required Privilege Level</b>	<b>interface</b> —To view this statement in the configuration. <b>interface-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Match Conditions In CoS Rules on page 557</a></li><li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li></ul>

## destination-prefix-list (Services IDS)

---

<b>Syntax</b>	<code>destination-prefix-list <i>list-name</i> &lt;except&gt;;</code>
<b>Hierarchy Level</b>	<code>[edit services ids <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>from</b>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2.
<b>Description</b>	Specify the destination prefix list for rule matching. You configure the prefix list by including the <b>prefix-list</b> statement at the <b>[edit policy-options]</b> hierarchy level.
<b>Options</b>	<b><i>list-name</i></b> —Destination prefix list.  <b>except</b> —(Optional) Exclude the specified prefix list from rule matching.
<b>Required Privilege Level</b>	<b>interface</b> —To view this statement in the configuration. <b>interface-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li><li>• <a href="#">Configuring Match Conditions in IDS Rules on page 386</a></li></ul>

## destination-prefix-list (Services NAT)

<b>Syntax</b>	<code>destination-prefix-list <i>list-name</i> &lt;except&gt;;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2.
<b>Description</b>	<p>Specify the destination prefix list for rule matching. You configure the prefix list by including the <b>prefix-list</b> statement at the [edit <b>policy-options</b>] hierarchy level.</p> <p>If the <a href="#">translation-type</a> statement in the <b>then</b> statement of the nat rule is set to <b>stateful-nat-64</b>, the destination prefix list for rule matching must be within the range specified by the <a href="#">destination-prefix</a> statement in the <b>then</b> statement.</p>
<b>Options</b>	<p><b>list-name</b>—Destination prefix list.</p> <p><b>except</b>—(Optional) Exclude the specified prefix list from rule matching.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> </ul>

## destination-prefix-list (Services Stateful Firewall)


<b>Syntax</b>	<code>destination-prefix-list <i>list-name</i> &lt;except&gt;;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> <a href="#">stateful-firewall</a> <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2.
<b>Description</b>	<p>Specify the destination prefix list for rule matching. You configure the prefix list by including the <b>prefix-list</b> statement at the [edit <b>policy-options</b>] hierarchy level.</p>
<b>Options</b>	<p><b>list-name</b>—Destination prefix list.</p> <p><b>except</b>—(Optional) Exclude the specified prefix list from rule matching.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Match Conditions in Stateful Firewall Rules on page 360</a></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> </ul>

## destined-port

---

<b>Syntax</b>	<code>destined-port <i>port id</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">port-forwarding</a> <i>map-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	Specify the port from where traffic has to be forwarded.
<b>Options</b>	<i>port id</i> —The destination port number from where traffic will be forwarded.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">port-forwarding on page 852</a></li><li>• <a href="#">translated-port on page 942</a></li></ul>

## deterministic-port-block-allocation

<b>Syntax</b>	deterministic-port-block-allocation { block-size <i>block-size</i> ; include-boundary-addresses; }
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">pool</a> <i>pool-name</i> port]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	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.
<b>Options</b>	<b>block-size</b> —Maximum number of ports that can be allocated to a user.
<div>  <p><b>NOTE:</b> When a <b>block-size</b> of 0 is specified, block size is calculated according to the formula: <math>(64512 * \text{Number of IP addresses in the NAT Pool}) / \text{Number of subscribers}</math> where</p> <ul style="list-style-type: none"> <li>64512 is derived from (65535 - 1023) because the regular port assignments start from 1024.</li> <li>Number of subscribers is derived from the <b>from</b> clause of the applicable NAT rule.</li> </ul> </div>	
<p><b>Default:</b> 256</p> <p><b>Range:</b> 0 through 32,000</p> <p><b>include-boundary-addresses</b>—(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.</p>	
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Configuring Deterministic Port Block Allocation on page 171</a></li> </ul>

## dh-group

---

<b>Syntax</b>	dh-group (group1   group2   group5  group14   group19   group20);
<b>Hierarchy Level</b>	[edit services ipsec-vpn ike <a href="#">proposal</a> <i>proposal-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the IKE Diffie-Hellman prime modulus group to use for performing the new Diffie-Hellman exchange.
<b>Options</b>	<p><b>group1</b>—768-bit.</p> <p><b>group2</b>—1024-bit.</p> <p><b>group5</b>—1536-bit.</p> <p><b>group14</b>—2048-bit.</p> <p><b>group19</b>—256-bit random Elliptic Curve Group.</p> <p><b>group20</b>—384-bit random Elliptic Curve Group.</p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IKE Proposals on page 435</a></li></ul>

## dial-options

<b>Syntax</b>	<pre>dial-options {   ipsec-interface-id <i>name</i>;   l2tp-interface-id <i>name</i>;   (shared   dedicated); }</pre>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>sp-fpc/pic/port unit logical-unit-number</i>],          [edit interfaces <i>si-fpc/pic/port unit logical-unit-number</i>],          [edit logical-systems <i>logical-system-name</i> interfaces <i>sp-fpc/pic/port unit logical-unit-number</i>],          [edit logical-systems <i>logical-system-name</i> interfaces <i>si-fpc/pic/port unit logical-unit-number</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.          The [edit ...si-...] hierarchy levels introduced in Junos OS Release 11.4.</p>
<b>Description</b>	Specify the options for configuring logical interfaces for group and user sessions in L2TP or IPsec dynamic endpoint tunneling.
<b>Options</b>	<p><b>dedicated</b>—(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><b>ipsec-interface-id <i>name</i></b>—(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><b>l2tp-interface-id <i>name</i></b>—Interface identifier that must be replicated at the [edit access profile <i>name</i>] hierarchy level.</p> <p><b>shared</b>—(LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.          interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 684</a></li> <li>• <a href="#">Configuring Dynamic Endpoints for IPsec Tunnels on page 495</a></li> <li>• <a href="#">Configuring Options for the LNS Inline Services Logical Interface</a></li> </ul>

## direction

---

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>rule</b> rule-name <b>term</b> term-name <b>then manual</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the direction in which manual SAs are applied.
<b>Options</b>	<p><b>bidirectional</b>—Apply the SA in both directions.</p> <p><b>inbound</b>—Apply the SA on inbound traffic.</p> <p><b>outbound</b>—Apply the SA on outbound traffic.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IPsec Rules on page 452</a></li></ul>

## dns-alg-pool

---

<b>Syntax</b>	<pre>dns-alg-pool dns-alg-pool;</pre>
<b>Hierarchy Level</b>	[edit <b>services</b> nat <b>rule</b> rule-name <b>term</b> term-name <b>then translated</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	Specify the Network Address Translation (NAT) pool for destination translation.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li></ul>



## dns-alg-prefix

<b>Syntax</b>	<code>dns-alg-prefix <i>dns-alg-prefix</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then translated</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	Set the Domain Name System (DNS) application-level gateway (ALG) 96-bit prefix for mapping IPv4 addresses to IPv6 addresses.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> </ul>

## drop-member-traffic (Aggregated Multiservices)

<b>Syntax</b>	<pre>drop-member-traffic {     <a href="#">rejoin-timeout</a> <i>rejoin-timeout</i>; }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> load-balancing-options member-failure-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	<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>
<b>Default</b>	If this statement is not configured, then the default behavior is to drop member traffic with a rejoin timeout of 120 seconds.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">member-failure-options (Aggregated Multiservices) on page 825</a></li> <li>• <a href="#">Understanding Aggregated Multiservices Interfaces on page 643</a></li> <li>• <a href="#">Example: Configuring an Aggregated Multiservices Interface (AMS) on page 652</a></li> </ul>

## ds-lite

<b>Syntax</b>	<pre>ds-lite ds-lite-software-concentrator {   auto-update-mtu;   copy-dscp;   flow-limit flow-limit   session-limit-per-prefix session-limit-per-prefix;   mtu-v6 mtu-v6;   software-address software-address; }</pre>
<b>Hierarchy Level</b>	[edit services software <a href="#">software-concentrator</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.4.</p> <p><b>auto-update-mtu</b> option introduced in Junos OS Release 10.4.</p> <p><b>copy-dscp</b> option introduced in Junos OS Release 11.2.</p> <p><b>mtu-v6</b> option introduced in Junos OS Release 10.4.</p> <p><b>software-address</b> option introduced in Junos OS Release 10.4.</p>
<b>Description</b>	Configure settings for a DS-Lite concentrator used to process IPv4 packets encapsulated in IPv6.
<b>Options</b>	<p><b>ds-lite-software-concentrator</b>—Name applied to a DS-Lite software concentrator.</p> <p><b>auto-update-mtu</b>—This option is not currently supported.</p> <p><b>copy-dscp</b>—Copy DSCP information to IPv4 headers during decapsulation.</p> <p><b>flow-limit</b>—Maximum number of IPv4 flows per software.</p> <p><b>Range:</b> 0 through 16384 flows</p> <p><b>mtu-v6</b>—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.</p> <p><b>Range:</b> 0 through 9192 bytes</p> <p><b>session-limit-per-prefix</b>—Maximum number of sessions per B4 subnet prefix. (0 through 16384).</p> <p><b>Range:</b> 0 through 16384 sessions</p> <p><b>software-address</b>—Address of the DS-Lite software concentrator.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a DS-Lite Software Concentrator on page 237</a></li> </ul>

## dscp

<b>Syntax</b>	<code>dscp (<i>alias</i>   <i>bits</i>);</code>
<b>Hierarchy Level</b>	[edit services cos <a href="#">application-profile</a> <i>profile-name</i> ftp data], [edit services cos <a href="#">application-profile</a> <i>profile-name</i> sip ( <a href="#">video</a>   <a href="#">voice</a> )], [edit services cos <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> then], [edit services cos <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> then ( <a href="#">reflexive</a>   <a href="#">reverse</a> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1.
<b>Description</b>	Define the Differentiated Services code point (DSCP) mapping that is applied to the packets.
<b>Options</b>	<i>alias</i> —Name assigned to a set of CoS markers.  <i>bits</i> —Mapping value in the packet header.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Actions in CoS Rules on page 558.</a></li> </ul>

## dynamic

<b>Syntax</b>	<pre>dynamic {   ike-policy <i>policy-name</i>;   ipsec-policy <i>policy-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> then]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define a dynamic IPsec SA.
<b>Options</b>	<p><a href="#">ike-policy</a> <i>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.</p> <p><a href="#">ipsec-policy</a> <i>policy-name</i>—Name of the IPsec policy. This statement is optional and the default policy is used if none is supplied.</p>
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Security Associations on page 415</a></li> </ul>

## ecmp-alb

---

<b>Syntax</b>	<pre>ecmp-alb {     apply-groups;     apply-groups-except;     tolerance; }</pre>
<b>Hierarchy Level</b>	[edit chassis]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.2.
<b>Description</b>	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.

---

<b>Options</b>	<p><b>apply-groups</b>—Specify the groups from which to inherit configuration data.</p> <p><b>apply-groups-except</b>—Specify the groups from which configuration data should not be inherited.</p> <p><b>tolerance</b>—Specify the adaptive tolerance in percentage.</p> <p><b>Default:</b> 20%.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

## ei-mapping-timeout

<b>Syntax</b>	<code>mapping-timeout <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit <a href="#">services</a> nat <a href="#">pool</a> <i>nat-pool-name</i>]</code>
<b>Release Information</b>	<code>ei-mapping-timeout</code> statement introduced in JUNOS Releases 12.3.
<b>Description</b>	Specify the duration for endpoint independent translations that use the specified NAT pool. This includes endpoint-independent mapping (EIM) and endpoint-independent filtering (EIF).
<b>Options</b>	<p><b><i>seconds</i></b>—Lifetime of endpoint independent mappings in seconds.</p> <p><b>Default:</b> 300</p> <p><b>Range:</b> 120 through 864,000</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Configuration Overview on page 51</a></li> </ul>

## eif-flow-limit

<b>Syntax</b>	<code>eif-flow-limit <i>number-of-flows</i></code>
<b>Hierarchy Level</b>	<code>[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then translated</a> <a href="#">secure-nat-mapping</a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3
<b>Description</b>	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 <code>eif-flow-limit</code> is configured as <i>n</i> , then only <i>n</i> inbound connections are allowed at a given instance of time. The <i>n</i> +1 and subsequent connections arriving when <i>n</i> connections are alive are dropped. A new inbound connection is allowed only when one of the <i>n</i> connections times out or is closed. This limit is applied for all type of flows.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Protecting CGN Devices Against Denial of Service (DOS) Attacks on page 251</a></li> </ul>

## enable-change-on-ams-redistribution

---

<b>Syntax</b>	enable-change-on-ams-redistribution;
<b>Hierarchy Level</b>	[edit services service-set <i>service-set-name</i> service-set-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	<p>Enable the service set to be bounced (reset) for dynamic NAT scenarios (dynamic NAT, NAT64, and NAT44) for service sets associated with aggregated multiservices (AMS) interfaces, when a member interface of an AMS bundle rejoins or a member interface failure occurs. When a member interface fails, the application resources (NAT pool in the case of dynamic NAT scenarios) and traffic load need to be rebalanced. For application resources to be rebalanced, which is the NAT pool for dynamic NAT environments, the NAT pool is split and allocated by the service PIC daemon (spd).</p> <p>By default, the bouncing of service sets and splitting of a NAT pool is disabled in dynamic NAT conditions, such as the translation type being dynamic-nat44, napt-44, and nat64, which enables backward compatibility with earlier Junos OS releases.</p> <p>To prevent the bouncing of a service set for resplitting the NAT pool and disabling the splitting of the NAT pool in dynamic NAT environments, you need not include the <b>enable-change-on-ams-redistribution</b> statement at the [edit services service-set <i>service-set-name</i> service-set-options] hierarchy level.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Enabling the Reset of Service Sets for Aggregated Multiservices Interfaces on page 13</a></li></ul>

## enable-rejoin (aggregated Multiservices)

<b>Syntax</b>	enable-rejoin;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> load-balancing-options member-failure-options redistribute-all-traffic]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	<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>
<b>Default</b>	If you do not configure this option, then the failed members do not automatically rejoin the <b>ams</b> interface even after coming back online.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">redistribute-all-traffic (Aggregated Multiservices) on page 861</a></li> <li>• <a href="#">Understanding Aggregated Multiservices Interfaces on page 643</a></li> <li>• <a href="#">Example: Configuring an Aggregated Multiservices Interface (AMS) on page 652</a></li> </ul>

## encapsulation

---

<b>Syntax</b>	<code>encapsulation type;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the logical link-layer encapsulation type.
<b>Options</b>	<b>atm-mlppp-llc</b> —For ATM2 IQ physical interfaces only, use Multilink Point-to-Point Protocol (MLPPP) over AAL5 LLC encapsulation.  <b>frame-relay-ppp</b> —For Frame Relay circuits, use Frame Relay PPP encapsulation.  <b>multilink-ppp</b> —By default, voice services logical interfaces use MLPPP encapsulation.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Encapsulation for Voice Services on page 669</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>



## encryption

<b>Syntax</b>	<pre> encryption {     algorithm <i>algorithm</i>;     key (ascii-text <i>key</i>   hexadecimal <i>key</i>); } </pre>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then manual</b> direction <i>direction</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>aes-128-cbc</b> , <b>aes-192-cbc</b> , and <b>aes-256-cbc</b> options added in Junos OS Release 7.6.
<b>Description</b>	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 415](#)

---

## encryption-algorithm

---

**Syntax** encryption-algorithm *algorithm*;

**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.

**Usage Guidelines** See *Configuring the Encryption Algorithm for an IKE Proposal* and *Configuring the Encryption Algorithm for an IPsec Proposal*.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

## establish-tunnels

<b>Syntax</b>	<code>establish-tunnels (immediately   on-traffic);</code>
<b>Hierarchy Level</b>	<code>[edit services ipsec-vpn]</code>
<b>Release Information</b>	Statement introduced in Release 8.5 of Junos OS.
<b>Description</b>	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.
<b>Options</b>	<ul style="list-style-type: none"> <li>• <b>immediately</b>—IKE is activated immediately after VPN configuration and configuration changes are committed.</li> <li>• <b>on-traffic</b>—IKE is activated only when data traffic flows and must to be negotiated.</li> </ul>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">IPsec Hierarchy Level on page 706</a></li> </ul>

## f-max-period

<b>Syntax</b>	<code>f-max-period <i>number</i>;</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i> <i>unit</i> <i>logical-unit-number</i> <i>compression</i> <i>rtp</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <i>unit</i> <i>logical-unit-number</i> <i>compression</i> <i>rtp</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the maximum number of compressed packets allowed between the transmission of full headers in a compressed Real-time Transport Protocol (RTP) traffic stream.
<b>Options</b>	<i>number</i> —Maximum number of packets. <b>Default:</b> 256
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Compression of Voice Traffic on page 667</a></li> </ul>

## facility-override (Service Sets)

---

<b>Syntax</b>	<code>facility-override <i>facility-name</i>;</code>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> <a href="#">syslog host</a> <i>hostname</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Override the default facility for system log reporting.
<b>Options</b>	<p><i>facility-name</i>—Name of the facility that overrides the default assignment. Valid entries are:</p> <ul style="list-style-type: none"><li><code>authorization</code></li><li><code>daemon</code></li><li><code>ftp</code></li><li><code>kernel</code></li><li><code>local0</code> through <code>local7</code></li><li><code>user</code></li></ul>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring System Logging for Service Sets on page 26</a></li></ul>

## facility-override (System Log Reporting)

<b>Syntax</b>	<code>facility-override <i>facility-name</i>;</code>
<b>Hierarchy Level</b>	[edit services l2tp <b>tunnel-group</b> <i>group-name</i> <b>syslog host</b> <i>hostname</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Override the default facility for system log reporting.
<b>Options</b>	<p><b><i>facility-name</i></b>—Name of the facility that overrides the default assignment. Valid entries include:</p> <ul style="list-style-type: none"> <li><b>authorization</b></li> <li><b>daemon</b></li> <li><b>ftp</b></li> <li><b>kernel</b></li> <li><b>local0</b> through <b>local7</b></li> <li><b>user</b></li> </ul>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring System Logging of L2TP Tunnel Activity on page 682</a></li> </ul>

## family (Aggregated Multiservices)

<b>Syntax</b>	<code>family <i>family</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>interface-unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	Configure protocol family information for the logical interface.
<b>Options</b>	<b><i>family</i></b> —Protocol family. Currently, only one option, <b>inet</b> (IP version 4 suite), is supported.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">unit (Aggregated Multiservices) on page 951</a></li> <li>• <a href="#">Understanding Aggregated Multiservices Interfaces on page 643</a></li> <li>• <a href="#">Example: Configuring an Aggregated Multiservices Interface (AMS) on page 652</a></li> </ul>

## family (Interfaces)

---

Syntax	<pre>family inet {     address address {         ...     }     service {         input {             [ service-set service-set-name &lt;service-filter filter-name&gt; ];             post-service-filter filter-name;         }         output {             [ service-set service-set-name &lt;service-filter filter-name&gt; ];         }     } }</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	<b>family</b> —Protocol family. Valid settings for service interfaces include <b>inet</b> (IPv4) and <b>mpls</b> .  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"><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces.</li><li>• <a href="#">Configuring the Address and Domain for Services Interfaces on page 24</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>

## family (Voice Services)

<b>Syntax</b>	<pre>family (inet   mlppp   ...) {     address address {         ...     }     bundle interface-name; }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure protocol family information for the logical interface.
<b>Options</b>	<p><i>family</i>—Protocol family:</p> <ul style="list-style-type: none"> <li>• <i>inet</i>—IP version 4</li> <li>• <i>mlppp</i>—MLPPP</li> </ul> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces.</li> <li>• <a href="#">Configuring Network Interfaces for Voice Services on page 670</a></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul>

## force-entry

---

<b>Syntax</b>	(force-entry   ignore-entry);
<b>Hierarchy Level</b>	[edit services ids <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>Specify handling of entries in the IDS events cache:</p> <ul style="list-style-type: none"><li>• <b>force-entry</b>—Ensure that the entry has a permanent place in the IDS cache after one event is registered.</li><li>• <b>ignore-entry</b>—Ensure that all IDS events are ignored.</li></ul>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Actions in IDS Rules on page 387</a></li></ul>

## forwarding-class (Services CoS)

---

<b>Syntax</b>	forwarding-class <i>class-name</i> ;
<b>Hierarchy Level</b>	[edit services cos <b>application-profile</b> <i>profile-name</i> <b>ftp data</b> ], [edit services cos <b>application-profile</b> <i>profile-name</i> <b>sip</b> ( <b>video</b>   <b>voice</b> )], [edit services cos <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> ], [edit services cos <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> ( <b>reflexive</b>   <b>reverse</b> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1.
<b>Description</b>	Define the forwarding class to which packets are assigned.
<b>Options</b>	<b>class-name</b> —Name of the target application.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Actions in CoS Rules on page 558</a>.</li></ul>



## forwarding-class (Services CoS Fragmentation Properties)

<b>Syntax</b>	forwarding-class <i>class-name</i> { (fragment-threshold <i>bytes</i>   no-fragmentation); multilink-class <i>number</i> ; }
<b>Hierarchy Level</b>	[edit class-of-service fragmentation-maps]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	For link services IQ (lsq) interfaces only, define a forwarding class name and associated fragmentation properties within a fragmentation map.  The <b>fragment-threshold</b> and <b>no-fragmentation</b> statements are mutually exclusive.
<b>Default</b>	If you do not include this statement, the traffic in forwarding class <i>class-name</i> is fragmented.
<b>Options</b>	<i>class-name</i> —Name of the forwarding class.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 570</a></li> </ul>

## fragment-limit

<b>Syntax</b>	fragment-limit <i>number-of-fragments</i> ;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> services-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Configure the maximum number of fragments permitted in a packet before the packet is dropped.
<b>Options</b>	<i>number-of-fragments</i> —Maximum number of fragments permitted. <b>Range:</b> 1 to 250 fragments. <b>Default:</b> 250 fragments.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Fragmentation Control for MS-DPC and MS-PIC Service Interfaces on page 30</a></li> </ul>

## fragment-threshold (Forwarding Class Maps)

---

<b>Syntax</b>	<code>fragment-threshold <i>bytes</i>;</code>
<b>Hierarchy Level</b>	[ <code>edit class-of-service fragmentation-maps <a href="#">forwarding-class</a> <i>class-name</i></code> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	For link services IQ ( <b>lsq</b> ) interfaces only, set the fragmentation threshold for an individual forwarding class.
<b>Default</b>	If you do not include this statement, the fragmentation threshold you set at the [ <code>edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> ] or [ <code>edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options</code> ] 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.
<b>Options</b>	<b>bytes</b> —Maximum size, in bytes, for multilink packet fragments. Any nonzero value must be a multiple of 64 bytes. <b>Range:</b> 128 through 16,320 bytes
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 570</a></li></ul>

## fragment-threshold (Interfaces LSQ)

<b>Syntax</b>	<code>fragment-threshold <i>bytes</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>lsq-fpc/pic/port</i> <b>unit</b> <i>logical-unit-number</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>lsq-fpc/pic/port</i> <b>unit</b> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	For voice services interfaces, set the fragmentation threshold, in bytes.
<b>Options</b>	<b>bytes</b> —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. <b>Range:</b> 128 through 16,320 bytes <b>Default:</b> 0 bytes (no fragmentation)
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Delay-Sensitive Packet Interleaving on page 668</a></li> </ul>

## fragmentation-map

<b>Syntax</b>	<code>fragmentation-map <i>map-name</i>;</code>
<b>Hierarchy Level</b>	[edit class-of-service interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	For link services IQ ( <b>lsq</b> ) interfaces only, associate a fragmentation map with a multilink PPP interface or MLFR FRF.16 DLCI.
<b>Default</b>	If you do not include this statement, traffic in all forwarding classes is fragmented.
<b>Options</b>	<b>map-name</b> —Name of the fragmentation map.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 570</a></li> </ul>

## fragmentation-maps

---

<b>Syntax</b>	<pre>fragmentation-maps {   map-name {     forwarding-class class-name {       (fragment-threshold bytes   no-fragmentation);       multilink-class number;     }   } }</pre>
<b>Hierarchy Level</b>	[edit class-of-service]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	For link services IQ (lsq) interfaces only, define fragmentation properties for individual forwarding classes.
<b>Default</b>	If you do not include this statement, traffic in all forwarding classes is fragmented.
<b>Options</b>	<p><i>map-name</i>—Name of the fragmentation map.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 570</a></li></ul>

## from (Services CoS)

---

<b>Syntax</b>	<pre> from {   application-sets set-name;   applications [ application-names ];   destination-address address;   destination-prefix-list list-name &lt;except&gt;;   source-address address;   source-prefix-list list-name &lt;except&gt;; } </pre>
<b>Hierarchy Level</b>	[edit services cos <b>rule</b> rule-name <b>term</b> term-name]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1.
<b>Description</b>	Specify input conditions for a CoS term.
<b>Options</b>	<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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS Rules on page 556</a></li> </ul>

## from (Services IDS)

---

<b>Syntax</b>	<pre>from {   application-sets set-name;   applications [ application-names ];   destination-address (address   any-unicast) &lt;except&gt;;   destination-address-range low minimum-value high maximum-value &lt;except&gt;;   source-address (address   any-unicast) &lt;except&gt;;   source-address-range low minimum-value high maximum-value &lt;except&gt;; }</pre>
<b>Hierarchy Level</b>	[edit services ids <a href="#">rule rule-name</a> <a href="#">term term-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify input conditions for the IDS term.
<b>Options</b>	<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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Match Conditions in IDS Rules on page 386</a></li></ul>

---

## from

---

<b>Syntax</b>	<pre> from {     destination-address address;     ipsec-inside-interface interface-name;     source-address address; } </pre>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify input conditions for the IPsec term.
<b>Options</b>	<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>
<b>Usage Guidelines</b>	See <i>Configuring Match Direction for IPsec Rules</i> .
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

## from (Services HCM)

---

<b>Syntax</b>	<pre> from {     url-list url-list-name;     url url_identifier {         host hostname;         request-url page-name;     } } </pre>
<b>Hierarchy Level</b>	[edit services hcm url-rule <i>url-rule-name</i> term <i>term-num</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Specify input conditions for the HCM term.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

## from (Services NAT)

---

<b>Syntax</b>	<pre>from {     application-sets set-name;     applications [ application-names ];     destination-address (address   any-unicast) &lt;except&gt;;     destination-address-range low minimum-value high maximum-value &lt;except&gt;;     source-address address (address   any-unicast) &lt;except&gt;;     source-address-range low minimum-value high maximum-value &lt;except&gt;; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> rule-name <a href="#">term</a> term-name]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify input conditions for the NAT term.
<b>Options</b>	<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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li></ul>



## from (Services Stateful Firewall)

<b>Syntax</b>	<pre> from {   application-sets set-name;   applications [ application-names ];   destination-address (address   any-unicast) &lt;except&gt;;   destination-address-range low minimum-value high maximum-value &lt;except&gt;;   destination-prefix-list list-name &lt;except&gt;;   source-address (address   any-unicast) &lt;except&gt;;   source-address-range low minimum-value high maximum-value &lt;except&gt;;   source-prefix-list list-name &lt;except&gt;; } </pre>
<b>Hierarchy Level</b>	[edit services stateful-firewall <a href="#">rule rule-name</a> <a href="#">term term-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify input conditions for a stateful firewall term.
<b>Options</b>	<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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Stateful Firewall Rules on page 359</a></li> </ul>

## ftp (Services CoS)

---

<b>Syntax</b>	<pre>ftp {   data {     dscp (<i>alias</i>   <i>bits</i>);     forwarding-class <i>class-name</i>;   } }</pre>
<b>Hierarchy Level</b>	[edit services cos <b>application-profile</b> <i>profile-name</i> ftp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3.
<b>Description</b>	Set the appropriate <b>dscp</b> and <b>forwarding-class</b> value for FTP.
<b>Default</b>	By default, the system does not alter the DSCP or forwarding class for FTP traffic.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Application Profiles</i></li><li>• <i>sip (Application Profile)</i></li></ul>

## hash-keys (Aggregated Multiservices)

<b>Syntax</b>	<pre>hash-keys {     egress-key (destination-ip   source-ip);     ingress-key (destination-ip   source-ip); }</pre>
<b>Hierarchy Level</b>	[edit services service-set <i>service-set-name</i> interface-service load-balancing-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	<p>Configure the hash keys used for load balancing in aggregated multiservices (AMS) for service applications (Network Address Translation [NAT], stateful firewall, application-level gateway [ALG], HTTP header enrichment, and mobility). The hash keys supported in the ingress and egress direction are the source IP address and destination IP address.</p> <p>Hash keys are used to define the load-balancing behavior among the various members in the AMS group. For example, if <b>hash-keys</b> is configured as <b>source-ip</b>, then the hashing would be performed based on the source IP address of the packet. Therefore, all packets with the same source IP address land on the same member. Hash keys must be configured with respect to the traffic direction: ingress or egress. For example, if <b>hash-keys</b> is configured as <b>source-ip</b> in the ingress direction, then it should be configured as <b>destination-ip</b> in the egress direction. This is required to ensure that the packets of the same flow reach the same member of the AMS group.</p> <p>The configuration of the ingress and egress hash keys is mandatory if you are using AMS for NAT. This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen. Refer to <a href="#">Table 24 on page 786</a> for the supported hash keys.</p> <p>The resource-triggered option enables anchor session PICs to use the load or resource information from the anchor services PICs to select the AMS member will anchor the services for the subscriber for load balancing among AMS members. In addition, for mobile subscriber-aware services (such as HTTP header enrichment), you must configure the <b>resource-triggered</b> statement, which means that the load balancing is not done using the ingress and egress keys.</p>

Table 24: Hash Keys Supported for AMS for Service Applications

Service Set at Ingress Interface			Service Set at Egress Interface	
Hash Keys for NAT				
NAT Type	Ingress hash key	Egress hash key	Ingress hash key	Egress hash key
source static	Destination IP address	Source IP address	Source IP address	Destination IP address
source dynamic	Source IP address	Destination IP address	Destination IP address	Source IP address
Network Address Port Translation (NAPT)	Source IP address	Destination IP address	Destination IP address	Source IP address
destination static	Source IP address	Destination IP address	Destination IP address	Source IP address
Hash Keys for Stateful Firewall				
Stateful Firewall	Destination IP address	Source IP address	Destination IP address	Source IP address
Stateful Firewall	Source IP address	Destination IP address	Source IP address	Destination IP address



**NOTE:** If NAT is used in the service set (along with stateful firewall and ALG), then the hash keys should be based on the NAT type; otherwise, the hash keys of the stateful firewall should be used.

## Options



**NOTE:** The `egress-keys` option is hidden and is deprecated in Junos OS Release 15.1 and later, and is only maintained for backward compatibility. It might be removed completely in a future software release. Load-balancing or steering of traffic occurs, based on the hash keys in the forward direction. Load-balancing of traffic also occurs, based on the hash keys in the reverse direction except in dynamic NAT scenarios (dynamic NAT, NAT64, and NAPT44). For interface-style services, the ingress hash-key is used for the forward direction and the egress hash-key is used for the reverse direction. These hash-keys are configured within the service-set definition by using the `ingress-key` and `egress-key` statements at the `[edit services service-set service-set-name interface-service load-balancing-options]` hierarchy level. For next-hop style services, the ingress hash-key on the inside-domain next-hop is used in the forward direction and the ingress hash-key (not the egress hash-key) on outside-domain next-hop is used for the reverse direction. These hash-keys are configured at the logical AMS interface level by using the `ingress-key` and `egress-key` statements at the `[edit interfaces amsN unit logical-unit-number load-balancing-options hash-keys]` hierarchy level.

**ingress-key destination-ip**—Use the destination IP address of the flow to compute the hash used in load balancing in the ingress flow direction.

**ingress-key source-ip**—Use the source IP address of the flow to compute the hash used in load balancing in the ingress flow direction.

**egress-key destination-ip**—Use the destination IP address of the flow to compute the hash used in load balancing. Configure the hash keys to be used in the egress flow direction. The configuration is mandatory if you are using AMS for Network Address Translation (NAT). This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen.

**egress-key source-ip**—Use the source IP address of the flow to compute the hash used in load balancing. Configure the hash keys to be used in the egress flow direction.

**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 on page 809](#)

## header-integrity-check

---

**Syntax**    header-integrity-check {  
              enable-all;  
              }

**Hierarchy Level**    [edit services service-set *service-set* [service-set-options](#)]

**Release Information**    Statement introduced in Release 13.2.

**Description**    Configure Junos OS to verify the packet header for anomalies in IP, TCP, UDP, and ICMP information and flag such anomalies and errors.



**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-tunnelin 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.

---

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.


**Related Documentation**    • [service-set-options on page 893](#)

## hello-interval

<b>Syntax</b>	<code>hello-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit services l2tp <i>tunnel-group name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the keepalive timer for L2TP tunnels.
<div>  <b>NOTE:</b> Subordinate statement support depends on the platform. See individual statement topics for more detailed support information. </div>	
<b>Options</b>	<p><b><i>seconds</i></b>—Interval, in seconds, after which the server sends a hello message if no messages are received. A value of <b>0</b> means that no hello messages are sent.</p> <p><b>Default:</b> 60 seconds</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Timers for L2TP Tunnels on page 681</a></li> <li>• <a href="#">Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces</a></li> </ul>

## hide-avps

---

<b>Syntax</b>	hide-avps;
<b>Hierarchy Level</b>	[edit services l2tp <a href="#">tunnel-group name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Hide L2TP attribute-value pairs if the secret shared between the two ends of the tunnel is known.
<div> <b>NOTE:</b> This statement is not supported for L2TP LNS on MX Series routers.</div>	
<b>Default</b>	Attribute-value pairs that can be hidden are exposed, even if the secret information is known.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Hiding Attribute-Value Pairs for L2TP Tunnels on page 682</a></li></ul>



## 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 809](#)
- [Understanding Aggregated Multiservices Interfaces on page 643](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 652](#)

## hint

---

<b>Syntax</b>	<code>hint [ <i>hint-strings</i> ];</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">pool</a> <i>nat-pool-name</i> <a href="#">pgcp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0.
<b>Description</b>	Configure a hint that enables the border gateway function (BGF) to choose a NAT pool by direction rather than by virtual interface. The BGF matches the configured hint with a termination hint located in the Direction field of a nonstandard termination ID.
<b>Default</b>	When no hint is configured, the BGF can choose any NAT pool associated with the virtual interface.
<b>Options</b>	<b><i>hint-string</i></b> —Alphanumeric string of up to three characters that the BGF uses to match with a termination hint located in the Direction field of a nonstandard termination ID. You can also include underscores (_) and hyphens (-) within the string. To specify a list of hints, use the format: [ <b><i>hint xx hint yy</i></b> ].
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## host (L2TP)

---

<b>Syntax</b>	<pre>host <i>hostname</i> {     services <i>severity-level</i>;     <a href="#">facility-override</a> <i>facility-name</i>;     <a href="#">log-prefix</a> <i>prefix-value</i>; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> l2tp <a href="#">tunnel-group</a> <i>group-name</i> <a href="#">syslog</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the hostname for the system logging utility.
<b>Options</b>	<b><i>hostname</i></b> —Name of the system logging utility host machine. This can be the local Routing Engine or an external server address.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring System Logging of L2TP Tunnel Activity on page 682</a></li></ul>

## host (service-set)

<b>Syntax</b>	<pre> host <i>hostname</i> {   class {     alg-logs;     ids-logs;     nat-logs;     packet-logs;     pcp-logs;     session-logs &lt;open   close&gt;;     stateful-firewall-logs ;   }   facility-override <i>facility-name</i>;   interface-service <i>prefix-value</i>;   log-prefix <i>prefix-value</i>   port <i>port-number</i>   services <i>severity-level</i>;   source-address <i>source-address</i> } </pre>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> <a href="#">syslog</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>class</b> option introduced in Junos OS Release 13.2.
<b>Description</b>	Specify the hostname for the system logging utility.
<b>Options</b>	<p><b>hostname</b>—Name of the system logging utility host machine.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring System Logging for Service Sets on page 26</a></li> </ul>

## host (Services HCM)

---

<b>Syntax</b>	host <i>hostname</i> ;
<b>Hierarchy Level</b>	[edit services hcm url-rule <i>url-rule-name</i> term <i>term-num</i> from url <i>url_identifier</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Specify a hostname for the matching URL for the <b>term</b> . A match for that term is considered when a URL matches the <b>hostname</b> and the <b>request-URL</b> within the same term.
<b>Options</b>	<b>hostname</b> —Name of the host for the URL rule.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## hot-standby

---

<b>Syntax</b>	hot-standby;
<b>Hierarchy Level</b>	[edit interfaces <i>rlsnumber</i> <a href="#">redundancy-options</a> ], [edit interfaces <i>rlsnumber:number</i> <a href="#">redundancy-options</a> ] [edit interfaces <i>rspnumber</i> <a href="#">redundancy-options</a> ] [edit interfaces <i>rmsnumber</i> <a href="#">redundancy-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 592</a></li><li>• <a href="#">Configuring AS or Multiservices PIC Redundancy on page 20</a></li></ul>

## icmp-code

---

<b>Syntax</b>	<code>icmp-code value;</code>
<b>Hierarchy Level</b>	[edit <a href="#">applications application application-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Internet Control Message Protocol (ICMP) code value.
<b>Options</b>	<b>value</b> —The ICMP code value. For a complete list, see “ <a href="#">Configuring the ICMP Code and Type</a> ” on page 329.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">ALG Descriptions on page 299</a></li> <li>• <a href="#">Configuring Application Sets on page 325</a></li> <li>• <a href="#">Configuring the ICMP Code and Type on page 329</a></li> <li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li> <li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li> </ul>

## icmp-type

---

<b>Syntax</b>	<code>icmp-type value;</code>
<b>Hierarchy Level</b>	[edit <a href="#">applications application application-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	ICMP packet type value.
<b>Options</b>	<b>value</b> —The ICMP type value, such as <b>echo</b> or <b>echo-reply</b> . For a complete list, see “ <a href="#">Configuring the ICMP Code and Type</a> ” on page 329.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">ALG Descriptions on page 299</a></li> <li>• <a href="#">Configuring Application Sets on page 325</a></li> <li>• <a href="#">Configuring the ICMP Code and Type on page 329</a></li> <li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li> <li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li> </ul>

## ids-rules

---

<b>Syntax</b>	(ids-rules <i>rule-name</i>   ids-rule-sets <i>rule-set-name</i> );
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Rules on page 14</a></li></ul>

## ignore-entry

---

See [force-entry](#)

## ike

```

Syntax  ike {
        proposal proposal-name {
            authentication-algorithm (sha1 | sha-256 | sha-384);
            authentication-method (pre-shared-keys | rsa-signatures);
            description description;
            dh-group (group14 | group19 | group20);
            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.



**NOTE:** In Junos FIPS mode, the aggressive option of the **mode** statement is not supported.

**Usage Guidelines** See *Configuring IKE Proposals* and *Configuring IKE Policies*.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

## ike-access-profile

---

<b>Syntax</b>	<code>ike-access-profile <i>profile-name</i>;</code>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> <a href="#">ipsec-vpn-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4.
<b>Description</b>	Define the access profile for the IPsec traffic on dynamic tunnels.
<b>Options</b>	<i>profile-name</i> —Identifier for access profile, which must match the name configured at the [edit access profile <i>name</i> client * ike] hierarchy level.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Dynamic Endpoints for IPsec Tunnels on page 495</a></li><li>• <a href="#">Configuring IPsec Service Sets on page 460</a></li></ul>

## inactivity-timeout

---

<b>Syntax</b>	<code>inactivity-timeout <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">applications</a> <a href="#">application</a> <i>application-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Inactivity timeout period, in seconds.
<b>Options</b>	<i>seconds</i> —Length of time the application is inactive before it times out. <b>Default:</b> 30 seconds
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">ALG Descriptions on page 299</a></li><li>• <a href="#">Configuring Application Sets on page 325</a></li><li>• <a href="#">Configuring the Inactivity Timeout Period on page 334</a></li><li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li><li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li></ul>



## initiate-dead-peer-detection

<b>Syntax</b>	<code>initiate-dead-peer-detection;</code>
<b>Hierarchy Level</b>	<code>[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2.
<b>Description</b>	Enable triggering of dead peer detection (DPD) hello messages to the remote peer for the specified tunnel.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPsec Rules on page 452</a></li> <li>• <a href="#">dead-peer-detection on page 742</a></li> <li>• <a href="#">backup-remote-gateway on page 730</a></li> </ul>

## input (Interfaces)

<b>Syntax</b>	<pre>input {   service-set <i>service-set-name</i> &lt;service-filter <i>filter-name</i>&gt;;   post-service-filter <i>filter-name</i>; }</pre>
<b>Hierarchy Level</b>	<code>[edit interface <i>interface-name</i> unit <i>logical-unit-number</i> family inet service],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define the input service sets and filters to be applied to traffic.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Applying Filters and Services to Interfaces on page 17</a></li> </ul>

## interface

---

<b>Syntax</b>	<code>interface <i>interface-name.unit-number</i>;</code>
<b>Hierarchy Level</b>	[edit services service-interface-pools pool <i>pool-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3.
<b>Description</b>	Add logical service interfaces to the pool of service interfaces.
<b>Options</b>	<p><b><i>interface-name.unit-number</i></b>—Name and logical unit number of the service interface.</p> <ul style="list-style-type: none"><li>• All interfaces in a pool must belong to the same service PIC or DPC.</li><li>• All interfaces assigned to the same service must be in the same pool.</li><li>• Logical interfaces cannot be in more than one pool.</li><li>• All interfaces must have either <b>family inet</b> or <b>family inet6</b> configured.</li><li>• Logical unit 0 cannot be configured in a service interface pool.</li><li>• You can configure up to 1000 logical interfaces in a service interface pool.</li></ul>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Session Border Control Solutions Guide Using BGF and IMSG</i></li></ul>

## interface-service

---

<b>Syntax</b>	<pre>interface-service {   service-interface <i>name</i>; }</pre>
<b>Hierarchy Level</b>	[edit services <b>service-set</b> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the device name for the interface service Physical Interface Card (PIC).
<b>Options</b>	<b>service-interface <i>name</i></b> —Name of the service device associated with the interface-wide service set.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Sets to be Applied to Services Interfaces on page 9</a></li></ul>

## 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 649](#)
- [Understanding Aggregated Multiservices Interfaces on page 643](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 652](#)

## interfaces (Voice Services)

---

<b>Syntax</b>	<code>interfaces { ... }</code>
<b>Hierarchy Level</b>	[edit]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure interfaces on the router.
<b>Default</b>	The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>

## interval

---

<b>Syntax</b>	<code>interval seconds;</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then dead-peer-detection]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	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 <b>interval</b> value is used for IKEv1 security associations (SAs) but not for IKEv2 SAs.)
<b>Options</b>	<b>seconds</b> —Number of seconds that the peer waits before sending a DPD request packet. <b>Range:</b> 1 through 180 seconds <b>Default:</b> 2 seconds
<b>Required Privilege Level</b>	security—To view this statement in the configuration. security-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IPsec Rules on page 452</a></li></ul>

## ipsec

```
Syntax  ipsec {
        proposal proposal-name {
            authentication-algorithm (hmac-sha-256);
            description description;
            encryption-algorithm algorithm;
            lifetime-seconds seconds;
            protocol (esp | bundle);
        }
        policy policy-name {
            description description;
            perfect-forward-secrecy {
                keys (group14 | group19 | group20);
            }
            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.

**Usage Guidelines** See *Configuring Security Associations*.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

## 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 452](#)
- [Configuring Dynamic Endpoints for IPsec Tunnels on page 495](#)

## ipsec-vpn-options

---

<b>Syntax</b>	<pre>ipsec-vpn-options {     anti-replay-window-size <i>bits</i>;     clear-dont-fragment-bit;     copy-dont-fragment-bit     ike-access-profile <i>profile-name</i>;     local-gateway <i>address</i>;     no-anti-replay;     passive-mode-tunneling;     set-dont-fragment-bit     trusted-ca [ <i>ca-profile-names</i> ];     tunnel-mtu <i>bytes</i>; }</pre>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify IP Security (IPsec) service options.
<b>Options</b>	The remaining statements are explained separately.
<b>Usage Guidelines</b>	See “ <a href="#">Configuring Service Rules</a> ” on page 14.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## ipsec-vpn-rules

---

<b>Syntax</b>	(ipsec-vpn-rules <i>rule-name</i>   ipsec-vpn-rule-sets <i>rule-set-name</i> );
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<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.
<b>Usage Guidelines</b>	See “ <a href="#">Configuring Service Rules</a> ” on page 14.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## ipv6-multicast-interfaces


<b>Syntax</b>	ipv6-multicast-interfaces (all   <i>interface-name</i> ) { disable; }
<b>Hierarchy Level</b>	[edit <a href="#">services nat</a> ], [edit services software]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1.
<b>Description</b>	Enable multicast filters on Ethernet interfaces when IPv6 NAT is used for neighbor discovery.
<b>Options</b>	<p><b>all</b>—Enable filters on all interfaces.</p> <p><b>disable</b>—Disable filters on the specified interfaces.</p> <p><b><i>interface-name</i></b>—Enable filters on a specific interface only.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPv6 Multicast Interfaces on page 238</a></li> </ul>

## l2tp-access-profile

<b>Syntax</b>	l2tp-access-profile <i>profile-name</i> ;
<b>Hierarchy Level</b>	[edit services l2tp <a href="#">tunnel-group name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the profile used to validate all L2TP connection requests to the local gateway address.
<b>Options</b>	<b><i>profile-name</i></b> —Identifier for the L2TP connection profile.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Access Profiles for L2TP Tunnel Groups on page 680</a></li> <li>• <a href="#">Configuring an L2TP Access Profile on the LNS</a></li> </ul>

## land-attack-check

---

<b>Syntax</b>	land-attack-check (ip-only   ip-port );
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> <a href="#">nat-options</a> ]
<b>Release Information</b>	Statement introduced with Junos OS Release 12.3.
<b>Description</b>	Enable land attack checks based on either IP address only or both IP address and IP port number.
<div> <b>NOTE:</b> If you do not configure this statement, there is no land attack check for hairpinning NAT packets.</div>	
<b>Options</b>	<p>ip-only—Land attack check is based on IP address only.</p> <p>ip-port—Land attack check is based on IP address and IP port number.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Rules on page 14</a></li><li>• <a href="#">max-sessions-per-subscriber on page 823</a></li></ul>

## learn-sip-register

---

<b>Syntax</b>	learn-sip-register;
<b>Hierarchy Level</b>	[edit <a href="#">applications</a> <a href="#">application</a> <i>application-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4.
<b>Description</b>	Activate SIP register to accept potential incoming SIP calls.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">ALG Descriptions on page 299</a></li><li>• <a href="#">Configuring Application Sets on page 325</a></li><li>• <a href="#">Configuring SIP on page 325</a></li><li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li><li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li></ul>



## lifetime-seconds

<b>Syntax</b>	<code>lifetime-seconds <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <i>ike proposal proposal-name</i> ], [edit services ipsec-vpn <i>ipsec proposal proposal-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the lifetime of an IKE or IPsec SA. This statement is optional.
<b>Options</b>	<i>seconds</i> —Lifetime <b>Default:</b> 3600 seconds (IKE); 28,800 seconds (IPsec) <b>Range:</b> 180 through 86,400
<b>Usage Guidelines</b>	See <i>Configuring the Lifetime for an IKE SA</i> and <i>Configuring the Lifetime for an IPsec SA</i> .
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

## link-layer-overhead

<b>Syntax</b>	<code>link-layer-overhead <i>percent</i>;</code>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	For link services IQ ( <i>lsq</i> ) 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.
<b>Options</b>	<i>percent</i> —Percentage of total bundle bandwidth to be set aside for link-layer overhead. <b>Range:</b> 0 through 50 percent <b>Default:</b> 0 percent
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS Scheduling Queues on Logical LSQ Interfaces on page 566</a></li> </ul>

## load-balance

---

<b>Syntax</b>	<pre>load-balance {     per-packet;     random; }</pre>
<b>Hierarchy Level</b>	[edit policy-options policy-statement <i>policy-name</i> then]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.2.
<b>Description</b>	Specify the type of load balancing of an equal-cost multipath (ECMP) in the forwarding table.
<b>Options</b>	<p><b>per-packet</b>—Load-balance on a per-packet basis.</p> <p><b>random</b>—Load-balance using packet random spray.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Junos OS Routing Protocols and Policies Configuration Guide for Security Devices</i></li></ul>

## 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 (Aggregated Multiservices) {
            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 801](#)
- [Understanding Aggregated Multiservices Interfaces on page 643](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 652](#)

---

## local-certificate

---

<b>Syntax</b>	<code>local-certificate <i>identifier</i>;</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>ike</b> <b>policy</b> <i>policy-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Name of the certificate that needs to be sent to the peer during the IKE authentication phase.
<b>Options</b>	<i>identifier</i> —Name of certificate.
<b>Usage Guidelines</b>	See <i>Configuring the Local Certificate for an IKE Policy</i> .
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

## local-gateway (IPSec)

<b>Syntax</b>	<code>local-gateway <i>address</i>;</code>
<b>Hierarchy Level</b>	<code>[edit services <a href="#">service-set</a> <i>service-set-name</i> <a href="#">ipsec-vpn-options</a>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define the local IPv4 or IPv6 address for the IPsec traffic.
<b>Options</b>	<i>address</i> —Local address.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Service Rules on page 14</a></li> </ul>

## local-gateway (L2TP LNS)

<b>Syntax</b>	<pre>local-gateway {   address <i>address</i>;   gateway-name <i>gateway-name</i>; }</pre>
<b>Hierarchy Level</b>	<code>[edit services l2tp <a href="#">tunnel-group</a> <i>name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the IP address or name for the local (LNS) gateway for L2TP tunnel.  The remaining statements are explained separately.
<b>Options</b>	<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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Local Gateway Address and PIC on page 680</a>.</li> <li>• <a href="#">Configuring L2TP Tunnel Groups on page 679</a></li> <li>• <a href="#">Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces</a></li> </ul>

## local-id

---

<b>Syntax</b>	<code>local-id (ipv4_addr <i>ipv4-address</i>   ipv6_addr <i>ipv6-address</i>   key-id <i>identifier</i>);</code>
<b>Hierarchy Level</b>	<code>[edit services ipsec-vpn ike policy <i>policy-name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <code>ipv6_addr</code> option added in Junos OS Release 7.6.
<b>Description</b>	Specify local identifiers for IKE Phase 1 negotiation. This statement is optional.
<b>Options</b>	<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.
<b>Required Privilege Level</b>	<code>system</code> —To view this statement in the configuration. <code>system-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Security Associations on page 415</a></li></ul>

## log-prefix (L2TP)

---

<b>Syntax</b>	<code>log-prefix <i>prefix-value</i>;</code>
<b>Hierarchy Level</b>	<code>[edit services l2tp tunnel-group <i>group-name</i> syslog host <i>hostname</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Set the system logging prefix value.
<b>Options</b>	<code><i>prefix-value</i></code> —System logging prefix value.
<b>Required Privilege Level</b>	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring System Logging of L2TP Tunnel Activity on page 682</a></li></ul>

## log-prefix (Services)

<b>Syntax</b>	<code>log-prefix <i>prefix-value</i>;</code>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> <a href="#">syslog host</a> <i>hostname</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Set the system logging prefix value.
<b>Options</b>	<i>prefix-value</i> —System logging prefix value.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring System Logging for Service Sets on page 26</a></li> </ul>

## logging (Services)

<b>Syntax</b>	<pre>logging {   <a href="#">traceoptions</a> {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable         no-world-readable&gt;;     flag <i>flag</i>;     no-remote-trace;   } }</pre>
<b>Hierarchy Level</b>	[edit services]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.0.
<b>Description</b>	Define global services properties.
<b>Options</b>	The remaining statement is explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Tracing Services PIC Operations on page 27</a></li> </ul>

## logging (Services IDS)

---

<b>Syntax</b>	<pre>logging {     syslog;     threshold rate; }</pre>
<b>Hierarchy Level</b>	[edit services ids <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Set logging values for this IDS term.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Actions in IDS Rules on page 387</a></li></ul>

## lsq-failure-options

---

<b>Syntax</b>	<pre>lsq-failure-options {     no-termination-request;     trigger-link-failure <i>interface-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>lsq-fpc/pic/port</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4.
<b>Description</b>	For link services IQ ( <b>lsq</b> ) interfaces only, define the failure recovery option settings.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Association between LSQ and SONET Interfaces on page 590</a></li></ul>



## 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 415](#)

## 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	<b>preferred-backup <i>preferred-backup</i></b> —Name of the preferred backup member interface. The member interface format is <b>mams-a/b/0</b> , where <b>a</b> is the Flexible PIC Concentrator (FPC) slot number and <b>b</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"><li>• <a href="#">high-availability-options (aggregated Multiservices) on page 791</a></li><li>• <a href="#">Understanding Aggregated Multiservices Interfaces on page 643</a></li><li>• <a href="#">Example: Configuring an Aggregated Multiservices Interface (AMS) on page 652</a></li></ul>

---

## mapping-refresh

---

<b>Syntax</b>	mapping-refresh (inbound   outbound   inbound-outbound);
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> then translated secure-nat-mapping]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3
<b>Description</b>	Specify how the flow timer should be refreshed based on the mapping refresh configured for all types of fwnat flows. For TCP flows, if <b>tcp-tickles</b> 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).
<b>Options</b>	<b>inbound</b> —Refresh the flow timer for inbound flows only.  <b>inbound-outbound</b> —Refresh the flow timer for all flows.  <b>outbound</b> —Refresh the flow timer for outbound flows only.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Protecting CGN Devices Against Denial of Service (DOS) Attacks on page 251</a></li></ul>

## mapping-timeout

---

<b>Syntax</b>	mapping-timeout <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">pool</a> <i>nat-pool-name</i> ]
<b>Release Information</b>	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.

---

<b>Description</b>	Specify the duration for mappings that use the specified NAT pool.
<b>Options</b>	<b>seconds</b> —Lifetime of mappings in seconds. <b>Default:</b> 300 <b>Range:</b> 120 through 864,000
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Source and Destination Addresses Network Address Translation Overview on page 52</a></li></ul>

## match-direction (Services CoS)

---

<b>Syntax</b>	match-direction (input   output   input-output);
<b>Hierarchy Level</b>	[edit <a href="#">services</a> cos <a href="#">rule</a> <i>rule-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1.
<b>Description</b>	Specify the direction in which the rule match is applied.
<b>Options</b>	<b>input</b> —Apply the rule match on the input side of the interface. <b>output</b> —Apply the rule match on the output side of the interface. <b>input-output</b> —Apply the rule match bidirectionally.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring CoS Rules</a></li></ul>

## match-direction (Services IDS)

---

<b>Syntax</b>	<code>match-direction (input   output   input-output);</code>
<b>Hierarchy Level</b>	[edit services ids <a href="#">rule rule-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the direction in which the rule match is applied.
<b>Options</b>	<p><b>input</b>—Apply the rule match on input.</p> <p><b>output</b>—Apply the rule match on output.</p> <p><b>input-output</b>—Apply the rule match bidirectionally.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Match Conditions in IDS Rules on page 386</a></li> </ul>

## match-direction

---

<b>Syntax</b>	<code>match-direction (input   output);</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <a href="#">rule rule-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the direction in which the rule match is applied.
<b>Options</b>	<p><b>input</b>—Apply the rule match on input.</p> <p><b>output</b>—Apply the rule match on output.</p>
<b>Usage Guidelines</b>	See <i>Configuring Match Direction for IPsec Rules</i> .
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

## match-direction (Services NAT)

---

<b>Syntax</b>	match-direction (input   output);
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the direction in which the rule match is applied.
<b>Options</b>	<b>input</b> —Apply the rule match on input. <b>output</b> —Apply the rule match on output.
<b>Required Privilege Level</b>	<b>interface</b> —To view this statement in the configuration. <b>interface-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li></ul>

## match-direction (Services Stateful Firewall)

---


<b>Syntax</b>	match-direction (input   output   input-output);
<b>Hierarchy Level</b>	[edit <a href="#">services</a> stateful-firewall <a href="#">rule</a> <i>rule-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the direction in which the rule match is applied.
<b>Options</b>	<b>input</b> —Apply the rule match on the input side of the interface. <b>output</b> —Apply the rule match on the output side of the interface. <b>input-output</b> —Apply the rule match bidirectionally.
<b>Required Privilege Level</b>	<b>interface</b> —To view this statement in the configuration. <b>interface-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Stateful Firewall Rules on page 359</a></li></ul>

## max-drop-flows

<b>Syntax</b>	<pre>max-drop-flows {     ingress <i>ingress-flows</i>;     egress <i>egress-flows</i>; }</pre>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3
<b>Description</b>	<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>
<b>Options</b>	<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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Service Set Limitations on page 15</a></li> </ul>


## max-flows

---

<b>Syntax</b>	<code>max-flows <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Maximum number of flows allowed for the service set.
<b>Options</b>	<i>number</i> —Maximum number of flows.
<hr/>	
<div> <b>NOTE:</b> 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/>	
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Set Limitations on page 15</a></li></ul>



## max-sessions-per-subscriber

<b>Syntax</b>	<code>max-sessions-per-subscriber session-number;</code>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> <a href="#">nat-options</a> ]
<b>Release Information</b>	Statement introduced with Junos OS Release 12.3.
<b>Description</b>	Set the maximum number of sessions from a single subscriber allowed for network address translation/port for IPv4 to IPv4 (napt-44). This statement does not apply to other types of NAT. The maximum number of sessions per subscriber is 32,000 sessions.
<div>  <b>NOTE:</b> If you do not configure this statement, there is no limit to the number of sessions a subscriber can have. </div>	
<b>Options</b>	<p><i>session-number</i> —Maximum number of sessions a single subscriber can establish for NATP-44.</p> <p><b>Range:</b> 1 through 32000</p> <p><b>Default:</b> None</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Service Rules on page 14</a></li> <li>• <a href="#">land-attack-check on page 806</a></li> </ul>

## maximum

<b>Syntax</b>	<code>maximum number;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> services-options session-limit]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6.
<b>Description</b>	Specify the maximum number of sessions allowed simultaneously.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

## maximum-contexts

---

Syntax	maximum-contexts <i>number</i> <force>;
Hierarchy Level	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>compression</b> <i>rtp</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>compression</b> <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	<b>number</b> —Maximum number of contexts.  <b>force</b> —(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.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring Compression of Voice Traffic on page 667</a></li></ul>

## maximum-send-window

---

Syntax	maximum-send-window <i>packets</i> ;
Hierarchy Level	[edit services l2tp <b>tunnel-group</b> <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.



**NOTE:** This statement is not supported for L2TP LNS on MX Series routers.

---

Options	<b>packets</b> —Maximum number of packets the send window can hold at one time. <b>Default:</b> 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"><li>• <a href="#">Configuring Window Size for L2TP Tunnels on page 681</a></li></ul>

## 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 25 on page 825 displays the behavior of the member interface after the failure of the first Multiservices PIC. Table 26 on page 826 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 25: 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 26: 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	<b>drop-member-traffic</b>	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	<b>redistribute-all-traffic</b>	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 809](#)
- [Understanding Aggregated Multiservices Interfaces on page 643](#)
- [Example: Configuring an Aggregated Multiservices Interface \(AMS\) on page 652](#)

## member-interface (Aggregated Multiservices)

<b>Syntax</b>	<code>member-interface <i>interface-name</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> load-balancing-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	<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>  <p><b>NOTE:</b> The member interfaces that you specify must be members of aggregated Multiservices interfaces (mams-).</p> </div>	
The remaining statements are explained separately.	
<b>Options</b>	<p><i>interface-name</i>—Name of the member interface. The member interface format is <b>mams-a/b/0</b>, where <b>a</b> is the Flexible PIC Concentrator (FPC) slot number and <b>b</b> is the PIC slot number.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Aggregated Multiservices Interfaces on page 643</a></li> <li>• <a href="#">Example: Configuring an Aggregated Multiservices Interface (AMS) on page 652</a></li> <li>• <a href="#">load-balancing-options (Aggregated Multiservices) on page 809</a></li> </ul>

## message-rate-limit

<b>Syntax</b>	<code>message-rate-limit <i>messages-per-second</i></code>
<b>Hierarchy Level</b>	<pre> interfaces <i>interface-name</i> {   services-options {     <b>cg</b><i>n-pic</i>;     disable-global-timeout-override;     ignore-errors &lt;alg&gt; &lt;tcp&gt;;     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 {       <b>maximum</b> <i>number</i>;       rate <i>new-sessions-per-second</i>;     }     session-timeout <i>seconds</i>;     syslog {     }   } }</pre>
<b>Release Information</b>	Statement introduced Junos OS Release 11.1.
<b>Description</b>	Maximum system log messages per second allowed from this interface.
<div>  <p><b>NOTE:</b> The <code>message-rate-limit</code> command can be configured only for physical service interfaces (<code>sp-x/x/x</code>) and not for redundancy services PIC interfaces (<code>rspx</code>).</p> </div>	
<b>Options</b>	<p><b><i>messages-per-second</i></b>—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><b>Range:</b> 0 through 2147483647</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Configuring System Logging for Service Sets on page 26</a></li> </ul>

---

## mlfr-uni-nni-bundles-inline

---

<b>Syntax</b>	mlfr-uni-nni-bundles-inline <i>number</i> ;
<b>Hierarchy Level</b>	[edit chassis fpc <i>number</i> pic <i>number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	Specify the number of inline multilink frame relay user-to-network interface (UNI) and network-to-network interface (NNI) bundles.
<b>Options</b>	<i>number</i> —Specify the number of inline multilink frame relay UNI NNI bundles. <b>Range:</b> 1 through 255
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Junos OS to Support the Link Services PIC</i></li><li>• <a href="#">Inline MLPPP for WAN Interfaces Overview on page 603</a></li><li>• <i>Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces</i></li><li>• <i>Example: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces</i></li></ul>

## mode

---

<b>Syntax</b>	<code>mode (aggressive   main);</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>ike policy</b> <i>policy-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define an IKE policy mode.
<b>Default</b>	main
<b>Options</b>	<b>aggressive</b> —Takes half the number of messages of main mode, has less negotiation power, and does not provide identity protection.



**NOTE:** In Junos FIPS mode, the **aggressive** option is not supported.

---

**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.

<b>Usage Guidelines</b>	See <i>Configuring the Mode for an IKE Policy</i> .
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

## mss

---

<b>Syntax</b>	<code>mss value;</code>
<b>Hierarchy Level</b>	[edit services ids <b>rule rule-name term term-name then syn-cookie</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the maximum segment size (MSS) value used in Transmission Control Protocol (TCP) delayed binding.
<b>Options</b>	<b>value</b> —MSS value. <b>Default:</b> 1500 <b>Range:</b> 128 through 8192
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Actions in IDS Rules on page 387</a></li></ul>



## multi-link-layer-2-inline

<b>Syntax</b>	<code>multi-link-layer-2-inline;</code>
<b>Hierarchy Level</b>	<code>[edit chassis fpc <i>number</i> pic <i>number</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	Enable inline Layer 2 bundling services.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Junos OS to Support the Link Services PIC</i></li> <li>• <a href="#">Inline MLPPP for WAN Interfaces Overview on page 603</a></li> <li>• <i>Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces</i></li> <li>• <i>Example: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces</i></li> </ul>

## multilink-class

<b>Syntax</b>	<code>multilink-class <i>number</i>;</code>
<b>Hierarchy Level</b>	<code>[edit class-of-service fragmentation-maps <i>map-name</i> forwarding-class <i>class-name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>For link services IQ (<b>lsq</b>) interfaces only, map a forwarding class into a multiclass MLPPP (MCML).</p> <p>The <b>multilink-class</b> statement and <b>no-fragmentation</b> statements are mutually exclusive.</p>
<b>Options</b>	<p><b><i>number</i></b>—The multilink class assigned to this forwarding class.</p> <p><b>Range:</b> 0 through 7</p> <p><b>Default:</b> None</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 570</a></li> <li>• <a href="#">Configuring Multiclass MLPPP on LSQ Interfaces on page 606</a></li> <li>• <i>Configuring Fragmentation by Forwarding Class</i></li> <li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li> <li>• <a href="#">multilink-max-classes on page 832</a></li> </ul>

## multilink-max-classes

---

<b>Syntax</b>	<code>multilink-max-classes <i>number</i>;</code>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	For link services IQ (lsq) interfaces only, configure the number of multilink classes to be negotiated when a link joins the bundle.
<b>Options</b>	<b><i>number</i></b> —The number of multilink classes to be negotiated when a link joins the bundle. <b>Range:</b> 1 through 8 <b>Default:</b> None
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Multiclass MLPPP on LSQ Interfaces on page 606</a></li></ul>

## nat-options

---

<b>Syntax</b>	<pre>nat-options {   <a href="#">land-attack-check</a> (ip-only   ip-port);   <a href="#">max-sessions-per-subscriber</a> <i>session-number</i>;   <a href="#">stateful-nat64</a> {     <a href="#">clear-dont-fragment-bit</a>;   } }</pre>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced with Junos OS Release 12.1. <a href="#">land-attack-check</a> and <a href="#">max-sessions-per-subscriber</a> statements added in 13.3.
<b>Description</b>	Specify parameters for NAT operation.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Rules on page 14</a></li><li>• <a href="#">clear-dont-fragment-bit on page 738</a></li><li>• <a href="#">land-attack-check on page 806</a></li><li>• <a href="#">max-sessions-per-subscriber on page 823</a></li><li>• <a href="#">stateful-nat64 on page 911</a></li></ul>


---

## nat-rules

---

<b>Syntax</b>	(nat-rules <i>rule-name</i>   nat-rule-sets <i>rule-set-name</i> );
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Rules on page 14</a></li></ul>

## next-hop-service

<b>Syntax</b>	<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 <i>interface-type</i>;   service-interface-pool <i>name</i>; } </pre>
<b>Hierarchy Level</b>	[edit services <b>service-set</b> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>service-interface-pool</b> option added in Junos OS Release 9.3.
<b>Description</b>	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.
<b>Options</b>	<p><b>inside-service-interface <i>interface-name.unit-number</i></b>—Name and logical unit number of the service interface associated with the service set applied inside the network.</p> <p><b>outside-service-interface <i>interface-name.unit-number</i></b>—Name and logical unit number of the service interface associated with the service set applied outside the network.</p> <p><b>outside-service-interface-type <i>interface-type</i></b>—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><b>service-interface-pool <i>name</i></b>—Name of the pool of logical interfaces configured at the [edit services <b>service-interface-pools</b> <i>pool 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><b>NOTE:</b> <b>service-interface-pool</b> is not applicable for IP reassembly configuration on L2TP.</p> </div>	
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Configuring Service Sets to be Applied to Services Interfaces on page 9</a></li> </ul>

## no-anti-replay

<b>Syntax</b>	no-anti-replay;
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Disable IPsec antireplay service, which occasionally causes interoperability issues for security associations.
<b>Usage Guidelines</b>	See <i>Configuring or Disabling IPsec Anti-Replay</i> .
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.

## no-anti-replay (Services Service Set)

<b>Syntax</b>	no-anti-replay;
<b>Hierarchy Level</b>	[edit services service-set <i>service-set-name</i> ipsec-vpn-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0.
<b>Description</b>	<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 <b>no-anti-reply</b> 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 <b>anti-replay-window-size</b> 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.

<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPsec Service Sets on page 460</a></li> <li>• <a href="#">Configuring IPsec Rules on page 452</a></li> </ul>

## no-fragmentation

---

<b>Syntax</b>	no-fragmentation;
<b>Hierarchy Level</b>	[edit class-of-service fragmentation-maps <a href="#">forwarding-class</a> <i>class-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>For link services IQ (<b>lsq</b>) 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>
<b>Default</b>	If you do not include this statement, the traffic in forwarding class <i>class-name</i> is fragmented.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 570</a></li></ul>

## no-ipsec-tunnel-in-traceroute

---

<b>Syntax</b>	no-ipsec-tunnel-in-traceroute;
<b>Hierarchy Level</b>	[edit services ipsec-vpn]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0.
<b>Description</b>	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.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Security Associations on page 415</a></li></ul>

## no-per-unit-scheduler

---

<b>Syntax</b>	no-per-unit-scheduler;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 11.4.
<b>Description</b>	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 <b>shared-scheduler</b> statement are mutually exclusive.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Oversubscribing Interface Bandwidth</a></li> </ul>

## no-termination-request

---

<b>Syntax</b>	no-termination-request;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ppp-options], [edit interfaces <i>lsq-fpc/pic/port</i> <i>lsq-failure-options</i> ]
<b>Release Information</b>	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.
<b>Description</b>	Inhibit PPP termination-request messages to the remote host if the primary circuit fails.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Association between LSQ and SONET Interfaces on page 590</a></li> </ul>

## no-translation

---

<b>Syntax</b>	no-translation;
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	Specify that traffic is not to be translated.
<b>Options</b>	none
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li></ul>

## output

---

<b>Syntax</b>	output { [ <a href="#">service-set</a> <i>service-set-name</i> < <a href="#">service-filter</a> <i>filter-name</i> > ]; }
<b>Hierarchy Level</b>	[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> <a href="#">unit</a> <i>logical-unit-number</i> family inet service]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define the output service sets and filters to be applied to traffic.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Applying Filters and Services to Interfaces on page 17</a></li></ul>



## overload-pool

---

<b>Syntax</b>	<code>overload-pool <i>overload-pool-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then translated</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	Specify an address pool that can be used if the source pool becomes exhausted.
<b>Options</b>	<i>overload-pool-name</i> —Name of the overload pool.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> </ul>


## overload-prefix

---

<b>Syntax</b>	<code>overload-prefix <i>overload-prefix</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then translated</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	Specify the prefix that can be used if the source pool becomes exhausted.
<b>Options</b>	<i>overload-prefix</i> —Prefix value.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> </ul>

## passive-mode-tunneling

---

<b>Syntax</b>	<code>passive-mode-tunneling;</code>
<b>Hierarchy Level</b>	<code>[edit services <a href="#">service-set</a> <i>service-set-name</i> ipsec-vpn-options]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0.
<b>Description</b>	<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> <div><p><b>NOTE:</b> The <code>header-integrity-check</code> 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 <code>header-integrity-check</code> statement and the <code>passive-mode-tunneling</code> 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 <code>[edit 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 <code>[edit 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></div>
<b>Required Privilege Level</b>	<code>admin</code> —To view this statement in the configuration. <code>admin-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IPsec Service Sets on page 460</a></li></ul>

---

## pba-interim-logging-interval

---

<b>Syntax</b>	<code>pba-interim-logging-interval seconds</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> services-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.3.
<b>Description</b>	Port block allocation (PBA) generates one syslog entry per set of ports allocated to a subscriber. These logs are UDP based and can be lost in the network, especially for long running flows. Interim logging resends the above logs at a configured interval for all active blocks that have traffic on at least one block. For the MS-MIC and MS-MPC, log messages are generated for sessions which have a port in a block, even if the block has no traffic.
<b>Options</b>	<b>seconds</b> —Interval, in seconds, for re-sending of session logs <b>Default:</b> 0—This indicates that interim logging is not used. <b>Range:</b> 1800 to 86,400seconds.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring NAT Session Logs on page 219</a></li></ul>

## per-unit-scheduler

<b>Syntax</b>	<code>per-unit-scheduler;</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i>]</code>
<b>Release Information</b>	<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>
<b>Description</b>	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:** To enable per-unit scheduling on MX104 routers, configure the `per-unit-scheduler` statement at each member physical interface level of a particular aggregated Ethernet interface as well as at that aggregated Ethernet interface level. On other routing platforms, it is enough if you include this statement at the aggregated Ethernet interface level.



**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.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs</i></li> <li>• <i>vlan-tagging</i></li> <li>• <i>flexible-vlan-tagging</i></li> <li>• <i>Example: Applying Scheduling and Shaping to VLANs</i></li> <li>• <i>Configuring Virtual LAN Queuing and Shaping on PTX Series Routers</i></li> </ul>

## perfect-forward-secrecy

<b>Syntax</b>	<pre>perfect-forward-secrecy {     keys (group1   group2  group5  group14); }</pre>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>ipsec policy</b> <i>policy-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define Perfect Forward Secrecy (PFS). Creates single-use keys. This statement is optional.
<b>Options</b>	<p><b>keys</b>—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"> <li>• <b>group1</b>—768-bit.</li> <li>• <b>group2</b>—1024-bit.</li> <li>• <b>group5</b>—1536-bit.</li> <li>• <b>group14</b>—2048-bit.</li> </ul>
<b>Usage Guidelines</b>	See <i>Configuring Perfect Forward Secrecy</i> .
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.

## pgcp

---

Syntax	<pre>pgcp {     hint [ hint-strings ];     ports-per-session ports;     remotely-controlled;     transport [ transport-protocols ]; }</pre>
Hierarchy Level	[edit <b>services</b> nat <b>pool</b> <i>nat-pool-name</i> ]
Release Information	Statement introduced in Junos OS Release 8.4. <b>remotely-controlled</b> and <b>ports-per-session</b> statements added in Junos OS Release 8.5. <b>hint</b> statement added in Junos OS Release 9.0.
Description	Specify that the NAT pool is used exclusively by the BGF.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## pgcp-rules

---

Syntax	<pre>(pgcp-rules <i>rule-name</i>   pgcp-rules-sets <i>rule-set-name</i>);</pre>
Hierarchy Level	[edit <b>services</b> <b>service-set</b> <i>service-set-name</i> ]
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	<b>rule-name</b> —Identifier for the collection of terms that constitute this rule.  <b>rule-set-name</b> —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"><li>• <a href="#">Configuring Service Sets to be Applied to Services Interfaces on page 9</a></li></ul>

## 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 439](#)

## policy (IPsec)

---

<b>Syntax</b>	<pre>policy <i>policy-name</i> {     <i>description</i> <i>description</i>;     perfect-forward-secrecy {         keys (group1   group2);     }     proposals [ <i>proposal-names</i> ]; }</pre>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <a href="#">ipsec</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define an IPsec policy.
<b>Options</b>	<p><i>policy-name</i>—IPsec policy name.</p> <p>The remaining statements are explained separately.</p>
<b>Usage Guidelines</b>	See <i>Configuring IPsec Policies</i> .
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>



## 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 53](#)

## pool (Service Interface)

---

<b>Syntax</b>	<code>pool <i>pool-name</i> {     <b>interface</b> <i>interface-name.unit-number</i>; }</code>
<b>Hierarchy Level</b>	[edit services <a href="#">service-interface-pools</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3.
<b>Description</b>	Configure a service interface pool for VPN aggregation for the BGF feature.
<b>Options</b>	<p><i>pool-name</i>—Name of the service interface pool.</p> <p>The remaining options are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Interface Pools on page 16</a></li></ul>

## 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.



**NOTE:** Starting with Junos OS Release 15.1, the **preserve-port** and **preserve-range** functionalities are supported on MX Series routers with MS-MPCs and MS-MICs.

**preserve-range**—Preserve privileged port range after translation.

**random-allocation**—Allocate ports within a specified range randomly.

Other options 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 Source and Destination Addresses Network Address Translation Overview on page 52](#)
- [Configuring Address Pools for Network Address Port Translation \(NAPT\) Overview on page 103](#)

## port (Services Voice)

<b>Syntax</b>	port { minimum <i>port-number</i> ; maximum <i>port-number</i> ; }
<b>Hierarchy Level</b>	[edit interfaces <i>lsq-fpc/pic/port</i> <b>unit</b> <i>logical-unit-number</i> <b>compression</b> <i>rtp</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>lsq-fpc/pic/port</i> <b>unit</b> <i>logical-unit-number</i> <b>compression</b> <i>rtp</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	For voice services interfaces only, specify a range of User Datagram Protocol (UDP) destination port numbers in which RTP compression takes place.
<b>Options</b>	<b>minimum <i>port-number</i></b> —Specify the minimum port number. <b>Range:</b> 0 through 65,535  <b>maximum <i>port-number</i></b> —Specify the maximum port number. <b>Range:</b> 0 through 65,535
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Compression of Voice Traffic on page 667</a></li> </ul>

## port (System Log Messages)

<b>Syntax</b>	port <i>port-number</i> ;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> services-options syslog host <i>hostname</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1.
<b>Description</b>	UDP port for system log messages on the host. The default port is 514.
<b>Options</b>	<b><i>port-number</i></b> —Port number for system log messages.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring System Logging for Services Interfaces</a></li> </ul>

## port-forwarding

---

<b>Syntax</b>	<code>port-forwarding <i>map-name</i> {     <i>destined-port</i>;     <i>translated-port</i>; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	Specify the mapping for port forwarding.
<b>Options</b>	<i>map-name</i> —Identifier for the port forwarding map.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Port Forwarding for Static Destination Address Translation on page 173</a></li><li>• <a href="#">Configuring Port Forwarding Without Destination Address Translation on page 176</a></li></ul>

## port-forwarding-mappings

---

<b>Syntax</b>	<code>port-forwarding-mappings <i>map-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	Specify the name for mapping port forwarding in a Network Address Translation configuration.
<b>Options</b>	<i>map-name</i> —Identifier for the port forwarding mapping.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Port Forwarding for Static Destination Address Translation on page 173</a></li><li>• <a href="#">Configuring Port Forwarding Without Destination Address Translation on page 176</a></li></ul>

## ports-per-session

---

<b>Syntax</b>	<code>ports-per-session <i>ports</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">pool</a> <i>nat-pool-name</i> <a href="#">pgcp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4.
<b>Description</b>	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.
<b>Options</b>	<i>number-of-ports</i> —Number of ports to enable: 2 or 4 for combined voice and video services. <b>Default:</b> 2
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## post-service-filter

---

<b>Syntax</b>	<code>post-service-filter <i>filter-name</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">family</a> inet <a href="#">service input</a> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">family</a> inet <a href="#">service input</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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 <b>post-service-filter</b> statement is not supported when the service interface is on an MS-MIC or MS-MPC.
<b>Options</b>	<i>filter-name</i> —Identifier for the post-service filter.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Applying Filters and Services to Interfaces on page 17</a></li> </ul>

## ppp-access-profile

---

<b>Syntax</b>	<code>ppp-access-profile <i>profile-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit services l2tp <a href="#">tunnel-group</a> <i>name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.

---

<b>Options</b>	<i>profile-name</i> —Identifier for the PPP profile.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Access Profiles for L2TP Tunnel Groups on page 680</a></li></ul>

## pre-shared-key (Services IKE)

---

<b>Syntax</b>	<code>pre-shared-key (ascii-text <i>key</i>   hexadecimal <i>key</i>);</code>
<b>Hierarchy Level</b>	<code>[edit services ike <a href="#">policy</a> <i>policy-name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define a preshared key for an IKE policy.
<b>Options</b>	<i>key</i> —Value of preshared key. The key can be one of the following: <ul style="list-style-type: none"><li>• <b>ascii-text</b>—ASCII text key.</li><li>• <b>hexadecimal</b>—Hexadecimal key.</li></ul>
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IKE Policies on page 439</a></li></ul>



## preserve-interface

<b>Syntax</b>	<code>preserve-interface;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> sonet-options aps]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	<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"> <li>• Channelized OC3 IQ PIC</li> <li>• Channelized OC12 IQ PIC</li> <li>• Channelized STM1 IQ PIC</li> </ul> <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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Link State Replication for Redundant Link PICs on page 595</a></li> </ul>

## primary (Adaptive Services Interfaces)

<b>Syntax</b>	<code>primary interface-name;</code>
<b>Hierarchy Level</b>	[edit interfaces (rsp0   rsp1) <b>redundancy-options</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the primary adaptive services interface.
<b>Options</b>	<i>interface-name</i> —The identifier for the AS or Multiservices PIC interface, which must be of the form <b>sp-fpc/pic/port</b> .
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring AS or Multiservices PIC Redundancy on page 20</a></li> </ul>

## primary (Link Services IQ PIC Interfaces)

---

<b>Syntax</b>	<code>primary interface-name;</code>
<b>Hierarchy Level</b>	[edit interfaces rlsqnumber <a href="#">redundancy-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	Specify the primary Link Services IQ PIC interface.
<b>Options</b>	<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> .
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 592</a></li></ul>

## proposal (Services IKE)

---

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <a href="#">ike</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define an IKE proposal for a dynamic SA.
<b>Options</b>	<i>proposal-name</i> —IKE proposal name.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IKE Proposals on page 435</a></li></ul>

## proposal (Services IPsec VPN)

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <a href="#">ipsec</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define an IPsec proposal for a dynamic SA.
<b>Options</b>	<p><i>proposal-name</i>—IPsec proposal name.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPsec Proposals on page 445</a></li> </ul>

## proposals

<b>Syntax</b>	proposals [ <i>proposal-names</i> ];
<b>Hierarchy Level</b>	[edit services ipsec-vpn <a href="#">ike policy</a> <i>policy-name</i> ], [edit services ipsec-vpn <a href="#">ipsec policy</a> <i>policy-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define a list of proposals to include in the IKE or IPsec policy.
<b>Options</b>	<i>proposal-names</i> —List of IKE or IPsec proposal names.
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IKE Proposals on page 435</a></li> <li>• <a href="#">Configuring IPsec Proposals on page 445</a></li> </ul>

## protocol (Applications)

---

<b>Syntax</b>	<code>protocol type;</code>
<b>Hierarchy Level</b>	[edit <a href="#">applications application application-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Networking protocol type or number.
<b>Options</b>	<b>type</b> —Networking protocol type. The following text values are supported:  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.

---

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">ALG Descriptions on page 299</a></li><li>• <a href="#">Configuring Application Sets on page 325</a></li><li>• <a href="#">Configuring Application Protocol Properties on page 325</a></li><li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li><li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li></ul>

## protocol (IPSec)

<b>Syntax</b>	<code>protocol (ah   esp   bundle);</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>ipsec proposal</b> <i>proposal-name</i> ], [edit services ipsec-vpn <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> manual direction <i>direction</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define an IPsec protocol for a dynamic or manual SA.
<b>Options</b>	<b>ah</b> —Authentication Header protocol.  <b>esp</b> —Encapsulating Security Payload protocol.  <b>bundle</b> —AH and ESP protocol.
<b>Required Privilege Level</b>	<b>admin</b> —To view this statement in the configuration. <b>admin-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Security Associations on page 415</a></li> </ul>

## ptsp-rules

<b>Syntax</b>	<code>(ptsp-rules <i>rule-name</i>   ptsp-rules-sets <i>rule-set-name</i>);</code>
<b>Hierarchy Level</b>	[edit services <b>service-set</b> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.2.
<b>Description</b>	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.
<b>Options</b>	<b>rule-name</b> —Identifier for the collection of terms that constitute this rule.  <b>rule-set-name</b> —Identifier for the set of rules to be included.
<b>Required Privilege Level</b>	<b>interface</b> —To view this statement in the configuration. <b>interface-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Service Sets to be Applied to Services Interfaces on page 9</a></li> </ul>

## queues

---

<b>Syntax</b>	<code>queues [ <i>queue-numbers</i> ];</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">compression</a> <a href="#">rtp</a> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">compression</a> <a href="#">rtp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	For voice services interfaces only, assign queue numbers on which RTP compression takes place.
<b>Options</b>	<code>queues <i>queue-numbers</i></code> —Assign one or more of the following queues: <b>q0</b> , <b>q1</b> , <b>q2</b> , and <b>q3</b> .
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Compression of Voice Traffic on page 667</a></li></ul>

## reassembly-timeout

---

<b>Syntax</b>	<code>reassembly-timeout <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> services-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	The maximum acceptable time, in seconds, from the receipt of the first and latest fragments in a packet. When the number is exceeded, the packet is dropped.
<b>Options</b>	<code><i>seconds</i></code> —Maximum seconds allowed. <b>Range:</b> 1 to 60 seconds. <b>Default:</b> 4 seconds.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Fragmentation Control for MS-DPC and MS-PIC Service Interfaces on page 30</a></li></ul>

## receive-window

<b>Syntax</b>	<code>receive-window <i>packets</i>;</code>
<b>Hierarchy Level</b>	<code>[edit services l2tp <i>tunnel-group name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.

<b>Options</b>	<i>packets</i> —Maximum number of packets the receive window can hold at one time. <b>Default:</b> 16
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Window Size for L2TP Tunnels on page 681</a></li> </ul>

## redistribute-all-traffic (Aggregated Multiservices)

<b>Syntax</b>	<code>redistribute-all-traffic {     <i>enable-rejoin</i>; }</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i> load-balancing-options member-failure-options]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	<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>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Aggregated Multiservices Interfaces on page 643</a></li> <li>• <a href="#">Example: Configuring an Aggregated Multiservices Interface (AMS) on page 652</a></li> <li>• <a href="#">member-failure-options (Aggregated Multiservices) on page 825</a></li> </ul>

## redundancy-options (Adaptive Services Interfaces)

---

<b>Syntax</b>	<pre>redundancy-options {     primary sp-fpc/pic/port;     secondary sp-fpc/pic/port;     hot-standby }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>rspnumber</i> ] [edit interfaces <i>rmsnumber</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the primary and secondary (backup) adaptive services interfaces.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring AS or Multiservices PIC Redundancy on page 20</a></li></ul>

## redundancy-options (Link Services IQ PIC Interfaces)

---

<b>Syntax</b>	<pre>redundancy-options {     (hot-standby   warm-standby);     primary lsq-fpc/pic/port;     secondary lsq-fpc/pic/port; }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>rlsqnumber</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	Specify the primary and secondary (backup) Link Services IQ PIC interfaces.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 592</a></li></ul>



## redundancy-options (MS-MIC, MS-MPC)

<b>Syntax</b>	<pre> redundancy-options {   redundancy-peer {     ipaddress <i>address</i>   }   routing-instance <b>redundancy-options</b> <i>name</i> } </pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.3.
<b>Description</b>	Specify the primary and secondary (backup) adaptive services PIC interfaces.



**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.

<b>Options</b>	<p><i>ipaddress</i>—Internal IP address of the remote redundant PIC.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Inter-Chassis High Availability for Stateful Firewall and NAPT44 (MS-MPC, MS-MIC) on page 263</a></li> </ul>

## (reflexive | reverse)

---

<b>Syntax</b>	<pre>(reflexive   reverse) {   application-profile profile-name;   dscp (alias   bits);   forwarding-class class-name;   syslog; }</pre>
<b>Hierarchy Level</b>	[edit services cos <b>rule</b> rule-name <b>term</b> term-name <b>then</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1.
<b>Description</b>	<p><b>reflexive</b>—Applies the equivalent opposing CoS action to flows in the opposite direction.</p> <p><b>reverse</b>—Allows you to define CoS behavior for flows in the reverse direction.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring CoS Rules</i></li><li>• <a href="#">Configuring Reflexive and Reverse CoS Rule Actions on page 560</a></li></ul>

## rejoin-timeout (Aggregated Multiservices)

<b>Syntax</b>	<code>rejoin-timeout <i>rejoin-timeout</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> load-balancing-options member-failure-options drop-member-traffic]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	<p>Configure the time by when failed members (members in the <b>DISCARD</b> state) should rejoin the aggregated Multiservices (AMS) interface automatically. All members that do not rejoin by the configured time are moved to the <b>INACTIVE</b> state and the traffic meant for each of the members is dropped.</p> <p>If multiple members fail around the same time, then they are held in the <b>DISCARD</b> state using a single timer. When the timer expires, all the failed members move to <b>INACTIVE</b> state at the same time.</p>
<b>Default</b>	If you do not configure a value, the default value of 120 seconds is used.
<b>Options</b>	<p><i>rejoin-timeout</i>—Time, in seconds, by which a failed member must rejoin.</p> <p><b>Default:</b> 120 seconds</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Aggregated Multiservices Interfaces on page 643</a></li> <li>• <a href="#">Example: Configuring an Aggregated Multiservices Interface (AMS) on page 652</a></li> <li>• <a href="#">drop-member-traffic (Aggregated Multiservices) on page 759</a></li> </ul>

## remote-gateway

<b>Syntax</b>	<code>remote-gateway <i>address</i>;</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define the remote address to which the IPsec traffic is directed.
<b>Options</b>	<i>address</i> —Remote IPv4 or IPv6 address.
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPsec Rules on page 452</a></li> </ul>

## remote-id

---

<b>Syntax</b>	<pre>remote-id {   any-remote-id;   ipv4_addr [ values ];   ipv6_addr [ values ];   key_id [ values ]; }</pre>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <i>ikepolicy policy-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>ipv6_addr</b> option added in Junos OS Release 7.6. <b>any-remote-id</b> option added in Junos OS Release 8.2.
<b>Description</b>	Define the remote identification values to which the IKE policy applies.
<b>Options</b>	<b>any-remote-id</b> —Allow any remote address to connect. This option is supported only in dynamic configurations and cannot be configured with specific values.  <b>ipv4_addr [ values ]</b> —Define one or more IPv4 address identification values.  <b>ipv6_addr [ values ]</b> —Define one or more IPv6 address identification values.  <b>key_id [ values ]</b> —Define one or more key identification values.  <b>fqdn fqdn</b> —Fully-qualified domain name.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IKE Policies on page 439</a></li></ul>

## remotely-controlled

---

<b>Syntax</b>	<pre>remotely-controlled;</pre>
<b>Hierarchy Level</b>	[edit <i>services</i> nat <i>pool nat-pool-name pgcp</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5.
<b>Description</b>	Configure the addresses and ports in a NAT pool to be remotely controlled by the gateway controller.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.


## request-url

---

<b>Syntax</b>	<code>request-url <i>page-name</i> ;</code>
<b>Hierarchy Level</b>	<code>[edit services hcm url-rule <i>url-rule-name</i> term <i>term-num</i> from url <i>url_identifier</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Specify a request-URL to match the <b>term</b> . A match for the term is considered when a URL matches any hostname and any request-URL within a term.
<b>Options</b>	<i>page-name</i> —Page name of the request URL.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## replicate-services (MS-MIC, MS-MPC)


---

<b>Syntax</b>	<pre>replicate-services {     replication-threshold <i>seconds</i>;     stateful-firewall;     nat; }</pre>
<b>Hierarchy Level</b>	[edit services service-set <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.3.
<b>Description</b>	Configure the services replication options for inter-chassis high availability on MS-MIC and MS-MPC multiservices PICs.
<b>Options</b>	<p><b>replication-threshold <i>seconds</i></b>—Specify the number of seconds for the replication-threshold. When a flow has been active for more than the number of seconds specified as a threshold, flow state information is replicated to the backup device.</p> <hr/> <div> <b>NOTE:</b> Make sure that the replication-threshold value is than the <i>open-timeout</i> (the timeout period for establishing a TCP connection).</div> <hr/> <p><b>Default:</b> 180. This value is also the minimum.</p> <p><b>nat</b>—Replicate NAPT44 state information.</p> <p><b>stateful-firewall</b>—Replicate stateful firewall information.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Inter-Chassis High Availability for Stateful Firewall and NAPT44 (MS-MPC, MS-MIC) on page 263</a></li></ul>

## respond-bad-spi (Services IKE Policy)

<b>Syntax</b>	<code>respond-bad-spi <i>max-responses</i>;</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <a href="#">ike policy</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3.
<b>Description</b>	Enable response to invalid IPsec Security Parameter Index (SPI) values. If the security associations (SAs) between two peers of an IPsec VPN become unsynchronized, the device resets the state of a peer so that the two peers are synchronized.
<b>Options</b>	<b>max-responses</b> —Number of times to respond to invalid SPI values per gateway. <b>Range:</b> 1 through 30 <b>Default:</b> 5
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IKE Policies on page 439</a></li> </ul>

## retransmit-interval (Services)

<b>Syntax</b>	<code>retransmit-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit services l2tp <a href="#">tunnel-group name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the maximum retransmit interval for L2TP tunnels.
<div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> This statement is not supported for L2TP LNS on MX Series routers.</p> </div> </div>	
<b>Options</b>	<b>seconds</b> —Interval, in seconds, after which the server retransmits data if no acknowledgment is received. <b>Default:</b> 30 seconds
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Timers for L2TP Tunnels on page 681</a></li> </ul>

## rpc-program-number

---

<b>Syntax</b>	<code>rpc-program-number <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">applications application application-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Remote procedure call (RPC) or Distributed Computing Environment (DCE) value.
<b>Options</b>	<i>number</i> —RPC or DCE program value. <b>Range:</b> 100,000 through 400,000
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">ALG Descriptions on page 299</a></li><li>• <a href="#">Configuring an RPC Program Number on page 342</a></li><li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li><li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li></ul>

## routing-engine-services

---

<b>Syntax</b>	<code>routing-engine-services;</code>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set service-set service-set-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	When configuring a Routing Engine-based captive portal service, specify the service set options to apply to a service set. The Routing Engine-based captive portal supports a walled garden as a firewall service filter only. The services interfaces on the Routing Engine are identified with an si- prefix (for example, si-1/1/0). The si- interface contains all redirect and rewrite traffic and services for the Routing Engine.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">HTTP Redirect Service Overview</a></li><li>• <a href="#">service-interface (Routing Engine Services)</a></li></ul>



## rtp

<b>Syntax</b>	<pre> rtp {     f-max-period <i>number</i>;     maximum-contexts <i>number</i> &lt;force&gt;;     port {         minimum <i>port-number</i>;         maximum <i>port-number</i>;     }     queues [ <i>queue-numbers</i> ]; } </pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">compression</a> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">compression</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the RTP properties for voice services traffic.  The remaining statements are described separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Compression of Voice Traffic on page 667</a></li> </ul>

## rule (Services CoS)

<b>Syntax</b>	<pre> rule <i>rule-name</i> {   match-direction (input   output   input-output);   term <i>term-name</i> {     from {       application-sets <i>set-name</i>;       applications [ <i>application-names</i> ];       destination-address <i>address</i>;       destination-prefix-list <i>list-name</i> &lt;except&gt;;       source-address <i>address</i>;       source-prefix-list <i>list-name</i> &lt;except&gt;;     }     then {       application-profile <i>profile-name</i>;       dscp (<i>alias</i>   <i>bits</i>);       forwarding-class <i>class-name</i>;       syslog;       (reflexive   reverse) {         application-profile <i>profile-name</i>;         dscp (<i>alias</i>   <i>bits</i>);         forwarding-class <i>class-name</i>;         syslog;       }     }   } } </pre>
<b>Hierarchy Level</b>	[edit services cos], [edit services cos <b>rule-set</b> <i>rule-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1.
<b>Description</b>	Specify the rule the router uses when applying this service.
<b>Options</b>	<p><b>rule-name</b>—Identifier for the collection of terms that constitute this rule.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS Rules on page 556</a></li> </ul>

## 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*]

<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the rule the router uses when applying this service.
<b>Options</b>	<i>rule-name</i> —Identifier for the collection of terms that constitute this rule.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IDS Rules on page 384</a></li></ul>

## rule

```
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.

**Usage Guidelines**    *See [Configuring Match Direction for IPsec Rules](#).*

**Required Privilege**    interface—To view this statement in the configuration.  
**Level**                interface-control—To add this statement to the configuration.

## 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 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 55](#)

## 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.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation** • [Configuring Stateful Firewall Rules on page 359](#)



## rule (Software)

<b>Syntax</b>	<pre>rule <i>rule-name</i> {     match-direction (input   output);     term <i>term-name</i> {         then {             (ds-lite <i>ds-lite-software-concentrator</i>   v6rd <i>v6rd-software-concentrator</i>);         }     } }</pre>
<b>Hierarchy Level</b>	[edit services software], [edit services software rule-set <i>rule-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	Configure a rule to apply a software concentrator for a flow.
<b>Options</b>	<p><b><i>rule-name</i></b>—Identifier for the collection of terms that constitute this rule.</p> <p><b>input</b>—Apply the rule match on the input side of the interface.</p> <p><b>output</b>—Apply the rule match on the output side of the interface.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Software Rules on page 229</a></li> </ul>

## rule-set (Services CoS)

<b>Syntax</b>	<pre>rule-set <i>rule-set-name</i> {     [ <i>rule</i> <i>rule-name</i> ]; }</pre>
<b>Hierarchy Level</b>	[edit services cos]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1.
<b>Description</b>	Specify the rule set the router uses when applying this service.
<b>Options</b>	<b><i>rule-set-name</i></b> —Identifier for the collection of rules that constitute this rule set.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS Rule Sets</a></li> </ul>

## rule-set (Services IDS)

---

<b>Syntax</b>	<code>rule-set <i>rule-set-name</i> {     [ <i>rule</i> <i>rule-names</i> ]; }</code>
<b>Hierarchy Level</b>	[edit services ids]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the rule set the router uses when applying this service.
<b>Options</b>	<i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IDS Rule Sets on page 392</a></li></ul>

## rule-set

---

<b>Syntax</b>	<code>rule-set <i>rule-set-name</i> {     [ <i>rule</i> <i>rule-names</i> ]; }</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the rule set the router uses when applying this service.
<b>Options</b>	<i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set.
<b>Usage Guidelines</b>	See <i>Configuring IPsec Rule Sets</i> .
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## rule-set (Services NAT)

---

<b>Syntax</b>	<code>rule-set <i>rule-set-name</i> {     [ <i>rule</i> <i>rule-names</i> ]; }</code>
<b>Hierarchy Level</b>	[edit <code>services nat</code> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the rule set the router uses when applying this service.
<b>Options</b>	<i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> </ul>

## rule-set (Services Stateful Firewall)

---

<b>Syntax</b>	<code>rule-set <i>rule-set-name</i> {     [ <i>rule</i> <i>rule-names</i> ]; }</code>
<b>Hierarchy Level</b>	[edit <code>services stateful-firewall</code> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the rule set the router uses when applying this service.
<b>Options</b>	<i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Stateful Firewall Rule Sets on page 363</a></li> </ul>

## rule-set (Softwire)

---

<b>Syntax</b>	<code>rule-set <i>rule-set-name</i> {     rule <i>rule-name</i>; }</code>
<b>Hierarchy Level</b>	[edit services softwire]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	Specify the rule set the router uses when applying this service.
<b>Options</b>	<i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Softwire Rules on page 229</a></li></ul>

## secondary (Adaptive Services Interfaces)

---

<b>Syntax</b>	<code>secondary <i>interface-name</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces (rsp0   rsp1) <a href="#">redundancy-options</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the secondary (backup) adaptive services interface.
<b>Options</b>	<i>interface-name</i> —The identifier for the adaptive services interface, which must be of the form <i>sp-fpc/pic/port</i> .
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring AS or Multiservices PIC Redundancy on page 20</a></li></ul>

## secondary (Link Services IQ PIC Interfaces)

<b>Syntax</b>	<code>secondary interface-name;</code>
<b>Hierarchy Level</b>	[edit interfaces rlsqnumber <a href="#">redundancy-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	Specify the secondary (backup) Link Services IQ PIC interface.
<b>Options</b>	<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> .
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 592</a></li> </ul>

## secure-nat-mapping

<b>Syntax</b>	<pre>secure-nat-mapping {   mapping-refresh (inbound   outbound   inbound-outbound);   eif-flow-limit number-of-flows' }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> rule-name <a href="#">term</a> term-name <a href="#">then translated</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3
<b>Description</b>	Specify configuration options that help prevent or minimize the effect of attempted denial of service (DOS) attacks for NAT operations.
<b>Options</b>	The statements are explained separately.  —
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Protecting CGN Devices Against Denial of Service (DOS) Attacks on page 251</a></li> </ul>

## 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 103](#)

## 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*—Software concentrator name whose software-address is used in creating PCP mappings. The PCP server address must be the same as the software-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 4294667 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 153](#)

## 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 17](#)

## service-domain

---

Syntax	service-domain (inside   outside);
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 <b>next-hop-service</b> statement at the [edit services service-set <i>service-set-name</i> ] hierarchy level, the interface domain must match that specified with the <b>inside-service-interface</b> and <b>outside-service-interface</b> statements.
Options	<b>inside</b> —Interface used within the network.  <b>outside</b> —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"><li>• <a href="#">Configuring the Address and Domain for Services Interfaces on page 24</a></li></ul>

## service-filter (Interfaces)


---

Syntax	service-filter <i>filter-name</i> ;
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 <b>service-set</b> statement without a <b>service-filter</b> definition, Junos OS assumes the match condition is true and selects the service set for processing automatically.
Options	<b>filter-name</b> —Identifies the filter to be applied in service processing. You can include special characters, such as a forward slash (/), colon (:), or a period (.).
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Applying Filters and Services to Interfaces on page 17</a></li><li>• <a href="#">Junos OS Services Interfaces Library for Routing Devices</a></li></ul>

## service-interface (Adaptive Services Interfaces)

<b>Syntax</b>	<code>service-interface <i>interface-name</i>;</code>
<b>Hierarchy Level</b>	[edit services service-set <i>service-set-name</i> interface-service]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the name for the adaptive services interface associated with an interface-wide service set.
<b>Options</b>	<b>interface-name</b> —Identifier of the service interface.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Service Sets to be Applied to Services Interfaces on page 9</a></li> </ul>

## service-interface (L2TP Processing)

<b>Syntax</b>	<code>service-interface <i>interface-name</i>;</code>
<b>Hierarchy Level</b>	[edit services l2tp <a href="#">tunnel-group name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <i>si-fpc/pic/port</i> option added in Junos OS Release 11.4.
<b>Description</b>	Specify the service interface responsible for handling L2TP processing.
<div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> 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.</p> </div> </div>	
<b>Options</b>	<b>interface-name</b> —Name of the service interface. The interface type depends on the line card as follows: <ul style="list-style-type: none"> <li>• <i>sp-fpc/pic/port</i>—On AS or Multiservices PICs on M7i, M10i, and M120 routers.</li> <li>• <i>si-fpc/pic/port</i>—On MPCs on MX Series routers.</li> </ul>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Local Gateway Address and PIC on page 680</a></li> <li>• <a href="#">Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces</a></li> </ul>

## service-interface-pools

---

<b>Syntax</b>	<pre>service-interface-pools {     pool <i>pool-name</i> {         interface <i>interface-name.unit-number</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit services]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3.
<b>Description</b>	Configure service interface pools used for VPN aggregation.
<b>Options</b>	The options are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Interface Pools on page 16</a></li></ul>

## service-set (Interfaces)

---

<b>Syntax</b>	<pre>service-set <i>service-set-name</i>;</pre>
<b>Hierarchy Level</b>	[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)]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<i>service-set-name</i> —Identifies the service set.
<b>Required Privilege Level</b>	System—To view this statement in the configuration. System-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Applying Filters and Services to Interfaces on page 17</a></li></ul>

## 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. You can include special characters, such as a forward slash (/), colon (:), or a period (.).

**Range:** Up to 64 alphanumeric characters.

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*

## service-set-options

<b>Syntax</b>	<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;   enable-change-on-ams-redistribution; } </pre>
<b>Hierarchy Level</b>	[edit services service-set]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.1.</p> <p>The <b>enable-asymmetric-traffic-processing</b> and the <b>support-uni-directional-traffic</b> options were added in Junos OS Release 11.2.</p> <p>The <b>routing-engine-services</b> option was added in Junos OS Release 15.1.</p> <p><b>enable-change-on-ams-redistribution</b> option added in Junos OS Release 15.1</p>
<b>Description</b>	Specify the service set options to apply to a service set.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Service Sets to be Applied to Services Interfaces on page 9</a></li> <li>• <a href="#">Configuring APPID Support for Unidirectional Traffic</a></li> <li>• <a href="#">Enabling the Reset of Service Sets for Aggregated Multiservices Interfaces on page 13</a></li> </ul>

## services (NAT)

<b>Syntax</b>	<pre> services nat { ... } </pre>
<b>Hierarchy Level</b>	[edit]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define the service rules to be applied to traffic.
<b>Options</b>	<b>nat</b> —Identifies the NAT set of rules statements.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

## 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 387](#)



## set-dont-fragment-bit (Services Set)

<b>Syntax</b>	set-dont-fragment-bit;
<b>Hierarchy Level</b>	[edit services service-set <i>service-set-name</i> ipsec-vpn-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	<p>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 <b>set-dont-fragment-bit</b> statement at the [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> 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.</p> <p>By default, this statement is disabled on MS-MICs and MS-MPCs (the DF bit value is not configured in the outer header by default).</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPsec Service Sets on page 460</a></li> <li>• <a href="#">Configuring IPsec Rules on page 452</a></li> </ul>

## set-dont-fragment-bit (Services IPsec VPN)

<b>Syntax</b>	set-dont-fragment-bit;
<b>Hierarchy Level</b>	[edit services ipsec-vpn <i>rule rule-name</i> term <i>term-name</i> then]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	<p>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 <b>set-dont-fragment-bit</b> statement at the [edit services service-set <i>service-set-name</i> 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.</p> <p>By default, this statement is disabled on MS-MICs and MS-MPCs (the DF bit value is not configured in the outer header by default).</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## [sip-call-hold-timeout](#)

---

<b>Syntax</b>	<code>sip-call-hold-timeout <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">applications application application-name</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4.
<b>Description</b>	Timeout period for SIP calls placed on hold, in seconds.
<b>Options</b>	<b><i>seconds</i></b> —Length of time the application holds a SIP call open before it times out. <b>Default:</b> 7200 seconds <b>Range:</b> 0 through 36,000 seconds (10 hours)
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">ALG Descriptions on page 299</a></li><li>• <a href="#">Configuring SIP on page 325</a></li><li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li><li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li></ul>

## sip

<b>Syntax</b>	<pre> sip {     video {         dscp (alias   bits);         forwarding-class class-name;     }     voice {         dscp (alias   bits);         forwarding-class class-name;     } } </pre>
<b>Hierarchy Level</b>	[edit services cos <a href="#">application-profile</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3.
<b>Description</b>	Set the appropriate <b>dscp</b> and <b>forwarding-class</b> value for SIP traffic.
<b>Default</b>	By default, the system will not alter the DSCP or forwarding class for SIP traffic.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS Rules on page 556</a></li> </ul>

## snmp-command

<b>Syntax</b>	snmp-command <i>command</i> ;
<b>Hierarchy Level</b>	[edit <a href="#">applications</a> <a href="#">application</a> <i>application-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	SNMP command format.
<b>Options</b>	<i>command</i> —Supported commands are SNMP <b>get</b> , <b>get-next</b> , <b>set</b> , and <b>trap</b> .
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">ALG Descriptions on page 299</a></li> <li>• <a href="#">Configuring an SNMP Command for Packet Matching on page 342</a></li> <li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li> <li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li> </ul>

## snmp-trap-thresholds

<b>Syntax</b>	<pre>snmp-trap-thresholds {     flows high <i>high-threshold</i>   low <i>low-threshold</i>;     nat-address-port high <i>high-threshold</i>   low <i>low-threshold</i>; }</pre>
<b>Hierarchy Level</b>	[edit services <b>service-set</b> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.1.
<b>Description</b>	Configures SNMP flow thresholds for all flows for a service set or flows for all NAT pools configured for a service set..
<b>Options</b>	<p>The remaining options are described separately.</p> <p><b>flows high <i>high-threshold</i></b>—Configure the upper limit for all flows on the service set. The limit is expressed as a percentage of <b>max-flows</b> configured for the service set. When the number of active flows exceeds this limit, an SNMP trap is set.</p> <p><b>Default:</b> 90 percent of <b>max-flows</b></p> <p><b>flows low <i>low-threshold</i></b>—Configure the lower limit for all flows on the service set . The limit is expressed as a percentage of <b>max-flows</b> configured for the service set. When the number of active flows falls below this limit, an SNMP trap is set.</p> <p><b>Default:</b> 70 percent of <b>max-flows</b></p> <p><b>nat-address-port high <i>high-threshold</i></b>—Configure the upper limit for flows for all NAT pools on the service set. The limit is expressed as a percentage of <b>max-flows</b> configured for the service set. When the number of active flows exceeds this limit, an SNMP trap is set.</p> <p><b>Default:</b> 90 percent of <b>max-flows</b></p> <p><b>nat-address-port low <i>low-threshold</i></b>—Configure the lower limit for flows. The limit is expressed as a percentage of <b>max-flows</b> configured for the service set. When the number of active flows falls below this limit, an SNMP trap is set.</p> <p><b>Default:</b> 80 percent of <b>max-flows</b></p>



**NOTE:** SNMP traps that are generated when you modify the threshold value for flows of NAT address pools in a service set (by using the `snmp-trap-thresholds nat-address-port (high high-threshold | low low-threshold)` statement) are not effective in the PIC. Only the initial threshold value that is set is effective on the PIC and subsequent changes to the threshold value are not reflected on the PIC. As a workaround, for the configuration changes under the [edit services nat pool *nat-pool-name*] hierarchy level, you must deactivate and activate the relevant service-set to enable the updated configuration to become effective. Otherwise, you must reboot the PIC for the updated threshold value of to take effect.



**NOTE:** Until Junos OS Release 14.1, when the NAT pool utilization exceeded the high threshold value configured, an SNMP trap was sent. However, a similar SNMP trap was not triggered when the NAT pool utilization fell below the configured lower limit or threshold. Because NMS systems are being used to monitor and set alarm for threshold values, the absence of an SNMP trap when the low threshold value was reached caused NMS to retain an active alarm in the alarms list. As a result, starting with Release 14.2R1, an SNMP trap is generated when the NAT pool utilization reaches the lower threshold, thereby causing the alarm in NMS to be reset.

<b>Required Privilege</b>	interface—To view this statement in the configuration.
<b>Level</b>	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Service Set Limitations on page 15</a></li> </ul>

## softwire-concentrator

<b>Syntax</b>	<pre> softwire-concentrator {   ds-lite ds-lite-softwire-concentrator {     auto-update-mtu;     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>address</i>;   }   v6rd v6rd-softwire-concentrator {     ipv4-prefix <i>ipv4-prefix</i>;     v6rd-prefix <i>ipv6-prefix</i>;     mtu-v4 <i>mtu-v4</i>;   } } </pre>
<b>Hierarchy Level</b>	[edit services softwire]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	Configure settings for a softwire concentrator.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege</b>	interface—To view this statement in the configuration.
<b>Level</b>	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a DS-Lite Softwire Concentrator on page 237</a></li> <li>• <a href="#">Configuring a 6rd Softwire Concentrator on page 255</a></li> </ul>

## software-options

---


<b>Syntax</b>	<pre>software-options {   dslite-ipv6-prefix-length <i>dslite-ipv6-prefix-length</i> ; }</pre>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	Specify the IPv6 prefix length associated with a subscriber's basic broadband bridging device that is subject to a limited number of sessions.
<b>Options</b>	<p><b><i>dslite-ipv6-prefix-length</i></b>—Subnet prefix representing the size of the subnet subject to session limitation.</p> <p><b>Values:</b> 56, 64, 96, 128</p> <p><b>Default:</b> 0—no limitation.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">DS-Lite Per Subnet Limitation Overview on page 252</a></li></ul>

## software-rules

---

<b>Syntax</b>	<pre>(software-rule <i>rule-name</i>   software-rule-sets <i>rule-set-name</i>);</pre>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	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.
<b>Options</b>	<p><b><i>rule-name</i></b>—Identifier for the collection of terms that constitute this rule.</p> <p><b><i>rule-set-name</i></b>—Identifier for the set of rules to be included.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Rules on page 14</a></li></ul>

## source-address (Service Sets)

<b>Syntax</b>	<code>source-address <i>source-address</i></code>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> <a href="#">syslog host</a> <i>hostname</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.1.
<b>Description</b>	Specify a source address to record in system log messages that are directed to a remote machine specified in the <i>hostname</i> statement.
<div>  <p><b>NOTE:</b> 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>	
<b>Options</b>	<i>source-address</i> —A valid IP address, which is recorded as the message source in messages sent to the remote machines specified in the <i>host hostname</i> statement
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring System Logging for Service Sets on page 26</a></li> <li>• <a href="#">host on page 793</a></li> <li>• <a href="#">service-set on page 891</a></li> </ul>

## source-address (Services CoS)

<b>Syntax</b>	<code>source-address <i>address</i>;</code>
<b>Hierarchy Level</b>	[edit services cos <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. <i>address</i> option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.
<b>Description</b>	Source address for rule matching.
<b>Options</b>	<i>address</i> —Source IPv4 or IPv6 address or prefix value.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Match Conditions in a CoS Rule</a></li> <li>• <a href="#">Configuring Match Conditions In CoS Rules on page 557</a></li> </ul>

## source-address (Services IDS)

---

<b>Syntax</b>	source-address ( <i>address</i>   any-unicast) <except>;
<b>Hierarchy Level</b>	[edit services ids <a href="#">rule rule-name</a> <a href="#">term term-name</a> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>address</b> option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.
<b>Description</b>	Specify the source address for rule matching.
<b>Options</b>	<b>address</b> —Source IPv4 or IPv6 address or prefix value. <b>any-unicast</b> —Any unicast packet. <b>except</b> —(Optional) Exempt the specified address, prefix, or unicast packets from rule matching.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Match Conditions in IDS Rules on page 386</a></li></ul>

## source-address

---

<b>Syntax</b>	source-address <i>address</i> ;
<b>Hierarchy Level</b>	[edit services ipsec-vpn <a href="#">rule rule-name</a> <a href="#">term term-name</a> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the source address for rule matching.
<b>Options</b>	<b>address</b> —Source IP address.
<b>Usage Guidelines</b>	See <i>Configuring Match Conditions in IPsec Rules</i> .
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.



## source-address (Services NAT)

<b>Syntax</b>	source-address ( <i>address</i>   any-unicast) <except>;
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>any-unicast</b> and <b>except</b> options introduced in Junos OS Release 7.6. <b>address</b> option enhanced to support IPv6 addresses in Junos OS Release 8.5.
<b>Description</b>	Specify the source address for rule matching.
<b>Options</b>	<b>address</b> —Source IPv4 or IPv6 address or prefix value.  <b>any-unicast</b> —Any unicast packet.  <b>except</b> —(Optional) Prevent the specified address or unicast packets from being translated.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> </ul>

## source-address (Services Stateful Firewall)

<b>Syntax</b>	source-address ( <i>address</i>   any-unicast) <except>;
<b>Hierarchy Level</b>	[edit <a href="#">services</a> stateful-firewall <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>any-unicast</b> and <b>except</b> options introduced in Junos OS Release 7.6. <b>address</b> option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.
<b>Description</b>	Source address for rule matching.
<b>Options</b>	<b>address</b> —Source IPv4 or IPv6 address or prefix value.  <b>any-unicast</b> —Any unicast packet.  <b>except</b> —(Optional) Exclude the specified address, prefix, or unicast packets from rule matching.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Match Conditions in Stateful Firewall Rules on page 360</a></li> </ul>

## source-address-range (Services IDS)

---

<b>Syntax</b>	source-address-range low <i>minimum-value</i> high <i>maximum-value</i> <except>;
<b>Hierarchy Level</b>	[edit services ids <a href="#">rule rule-name term term-name from</a> ]
<b>Release Information</b>	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.
<b>Description</b>	Specify the source address range for rule matching.
<b>Options</b>	<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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Match Conditions in IDS Rules on page 386</a></li></ul>

## source-address-range (Services NAT)

---

<b>Syntax</b>	source-address-range low <i>minimum-value</i> high <i>maximum-value</i> <except>;
<b>Hierarchy Level</b>	[edit <a href="#">services nat rule rule-name term term-name from</a> ]
<b>Release Information</b>	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.
<b>Description</b>	Specify the source address range for rule matching.
<b>Options</b>	<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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li></ul>

## source-address-range (Services Stateful Firewall)

<b>Syntax</b>	source-address-range low <i>minimum-value</i> high <i>maximum-value</i> <except>;
<b>Hierarchy Level</b>	[edit services stateful-firewall <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>from</b> ]
<b>Release Information</b>	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.
<b>Description</b>	Source address range for rule matching.
<b>Options</b>	<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. <b>except</b> —(Optional) Exclude the specified address, prefix, or unicast packets from rule matching.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Match Conditions in Stateful Firewall Rules on page 360</a></li> </ul>

## source-pool

<b>Syntax</b>	source-pool <i>nat-pool-name</i> ;
<b>Hierarchy Level</b>	[edit <b>services</b> nat <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then translated</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the source address pool for translated traffic.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> </ul>

## source-port

---

<b>Syntax</b>	<code>source-port <i>port-number</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">applications application application-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Source port identifier.
<b>Options</b>	<i>port-value</i> —Identifier for the port. For a complete list, see “ <a href="#">Configuring Source and Destination Ports</a> ” on page 331.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">ALG Descriptions</a> on page 299</li><li>• <a href="#">Configuring Application Protocol Properties</a> on page 325</li><li>• <a href="#">Configuring Source and Destination Ports</a> on page 331</li><li>• <a href="#">Verifying the Output of ALG Sessions</a> on page 344</li></ul>

## source-prefix (Services IDS)

---

<b>Syntax</b>	<code>source-prefix <i>prefix-value</i>;</code>
<b>Hierarchy Level</b>	[edit services ids <a href="#">rule rule-name term term-name then aggregation</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the prefix value for source IPv4 address aggregation.
<b>Options</b>	<i>prefix-value</i> —Integer value. <b>Range:</b> 1 through 32
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Actions in IDS Rules</a> on page 387</li></ul>

## source-prefix (Services NAT)

<b>Syntax</b>	<code>source-prefix <i>source-prefix</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then translated</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. <i>source-prefix</i> option enhanced to support IPv6 addresses in Junos OS Release 8.5.
<b>Description</b>	Specify the source prefix for translated traffic.
<b>Options</b>	<i>source-prefix</i> —IPv4 or IPv6 source prefix value.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> </ul>

## source-prefix-ipv6

<b>Syntax</b>	<code>source-prefix-ipv6 <i>prefix-value</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> ids <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then aggregation</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5.
<b>Description</b>	Specify the prefix value for source IPv6 address aggregation.
<b>Options</b>	<i>prefix-value</i> —Integer value. <b>Range:</b> 1 through 128
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Actions in IDS Rules on page 387</a></li> </ul>

## source-prefix-list (Services CoS)

---

<b>Syntax</b>	source-prefix-list <i>list-name</i> <except>;
<b>Hierarchy Level</b>	[edit services cos <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>from</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2.
<b>Description</b>	Specify the source prefix list for rule matching. You configure the prefix list by including the <b>prefix-list</b> statement at the [edit policy-options] hierarchy level.
<b>Options</b>	<i>list-name</i> —Destination prefix list.  <b>except</b> —(Optional) Exclude the specified prefix list from rule matching.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring CoS Rules on page 556</a></li><li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li></ul>

## source-prefix-list (Services IDS)

---

<b>Syntax</b>	source-prefix-list <i>list-name</i> <except>;
<b>Hierarchy Level</b>	[edit services ids <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>from</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2.
<b>Description</b>	Specify the source prefix list for rule matching. You configure the prefix list by including the <b>prefix-list</b> statement at the [edit policy-options] hierarchy level.
<b>Options</b>	<i>list-name</i> —Destination prefix list.  <b>except</b> —(Optional) Exclude the specified prefix list from rule matching.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Match Conditions in IDS Rules on page 386</a></li><li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li></ul>

## source-prefix-list (Services NAT)


<b>Syntax</b>	source-prefix-list <i>list-name</i> <except>;
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2.
<b>Description</b>	Specify the source prefix list for rule matching. You configure the prefix list by including the <b>prefix-list</b> statement at the [edit <b>policy-options</b> ] hierarchy level.
<b>Options</b>	<p><b>list-name</b>—Destination prefix list.</p> <p><b>except</b>—(Optional) Exclude the specified prefix list from rule matching.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> </ul>

## source-prefix-list (Services Stateful Firewall)

<b>Syntax</b>	source-prefix-list <i>list-name</i> <except>;
<b>Hierarchy Level</b>	[edit <a href="#">services</a> stateful-firewall <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">from</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2.
<b>Description</b>	Specify the source prefix list for rule matching. You configure the prefix list by including the <b>prefix-list</b> statement at the [edit <b>policy-options</b> ] hierarchy level.
<b>Options</b>	<p><b>list-name</b>—Destination prefix list.</p> <p><b>except</b>—(Optional) Exclude the specified prefix list from rule matching.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Match Conditions in Stateful Firewall Rules on page 360</a></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> </ul>

## spi

---

<b>Syntax</b>	<code>spi spi-value;</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> manual direction <i>direction</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the SPI for an SA.
<b>Options</b>	<b>spi-value</b> —An arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet). <b>Range:</b> 256 through 16,639
<hr/>	
<div> <b>NOTE:</b> Use the auxiliary SPI when you configure the protocol statement to use the <b>bundle</b> option.</div> <hr/>	
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Security Associations on page 415</a></li></ul>

## stateful-firewall-rules

---

<b>Syntax</b>	(stateful-firewall-rules <i>rule-names</i>   stateful-firewall-rule-sets <i>rule-set-name</i> );
<b>Hierarchy Level</b>	[edit services <b>service-set</b> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<b>rule-name</b> —Identifier for the collection of terms that make up this rule. <b>rule-set-name</b> —Identifier for the set of rules to be included.
<b>Required Privilege Level</b>	System—To view this statement in the configuration. System-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Rules on page 14</a></li></ul>



## stateful-nat64

<b>Syntax</b>	stateful-nat64 { clear-dont-fragment-bit; }
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> <a href="#">nat-options</a> ]
<b>Release Information</b>	Statement introduced with Junos OS Release 12.1.
<b>Description</b>	Set parameters for stateful NAT64 operation.



**NOTE:** These parameters do not change the operation of other types of NAT.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Service Rules on page 14</a></li> <li>• <a href="#">clear-dont-fragment-bit on page 738</a></li> </ul>

## syslog (Services CoS)

<b>Syntax</b>	syslog;
<b>Hierarchy Level</b>	[edit services cos <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then</a> ], [edit services cos <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then</a> ( <a href="#">reflexive</a>   <a href="#">reverse</a> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1.
<b>Description</b>	Enable system logging. The system log information from the Adaptive Services or Multiservices PIC is passed to the kernel for logging in the <b>/var/log</b> directory. This setting overrides any <b>syslog</b> statement setting included in the service set or interface default configuration.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Actions in a CoS Rule</a></li> <li>• <a href="#">Configuring Actions in CoS Rules on page 558</a></li> </ul>

## syslog (Services IDS)

---

<b>Syntax</b>	syslog;
<b>Hierarchy Level</b>	[edit services ids <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> <b>logging</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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 <code>/var/log</code> directory.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Actions in IDS Rules on page 387</a></li></ul>

## syslog

---

<b>Syntax</b>	syslog;
<b>Hierarchy Level</b>	[edit services ipsec-vpn <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Usage Guidelines</b>	See <i>Configuring Actions in IPsec Rules</i> .
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## syslog (Services L2TP)

**Syntax** `syslog {  
     host hostname {  
         services severity-level;  
         facility-override facility-name;  
         log-prefix prefix-value;  
     }  
}`

**Hierarchy Level** [edit **services** l2tp **tunnel-group** *group-name*]

**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 `/var/log/l2tpd` 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**

- [Configuring System Logging of L2TP Tunnel Activity on page 682](#)

## syslog (Services NAT)

**Syntax** `syslog;`

**Hierarchy Level** [edit **services** nat **rule** *rule-name* **term** *term-name* **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**

- [Network Address Translation Rules Overview on page 55](#)

## syslog (Services Service Set)

---

Syntax	<pre>syslog {   host hostname {     class {       alg-logs;       ids-logs;       nat-logs;       packet-logs;       pcp-logs;       session-logs &lt;open   close&gt;;       stateful-firewall-logs ;     }     services severity-level;     facility-override facility-name;     interface-service prefix-value;   } }</pre>
Hierarchy Level	[edit services <a href="#">service-set</a> <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 <i>Configuring System Logging for Services Interfaces</i> .
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"><li>• <a href="#">Configuring System Logging for Service Sets on page 26</a></li></ul>

---

## syslog (Services Stateful Firewall)

---

<b>Syntax</b>	syslog;
<b>Hierarchy Level</b>	[edit services stateful-firewall <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Enable system logging. The system log information from the Adaptive Services or Multiservices PIC is passed to the kernel for logging in the <b>/var/log</b> directory. This setting overrides any <b>syslog</b> statement setting included in the service set or interface default configuration.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Actions in Stateful Firewall Rules on page 362</a></li></ul>

## syn-cookie

---

<b>Syntax</b>	<pre>syn-cookie {     mss value;     threshold rate; }</pre>
<b>Hierarchy Level</b>	[edit services ids <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<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>
<b>Options</b>	The remaining statements are described separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Actions in IDS Rules on page 387</a></li></ul>

---

## tcp-mss

---

<b>Syntax</b>	<code>tcp-mss <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5.
<b>Description</b>	Specify the TCP Maximum Segment Size (MSS) allowed for the service set.
<b>Options</b>	<i>number</i> —MSS value.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Service Set Limitations on page 15</a></li></ul>

## 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 556.](#)



## 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 384](#)  
**Documentation**

## term

```
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-sha-256);
                        key (ascii-text key | hexadecimal key);
                    }
                    auxiliary-spi spi-value;
                    encryption {
                        algorithm algorithm;
                        key (ascii-text key | hexadecimal key);
                    }
                    protocol (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.

**Usage Guidelines** See *Configuring Match Direction for IPsec Rules*.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

## 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 55](#)

## 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.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**    • [Configuring Stateful Firewall Rules on page 359](#)

## then (Services CoS)

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit services cos <b>rule</b> rule-name <b>term</b> term-name]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1.
<b>Description</b>	<p>Define the CoS term actions.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Actions in a CoS Rule</i></li> <li>• <a href="#">Configuring Actions in CoS Rules on page 558</a></li> </ul>

## then (Services HCM)

<b>Syntax</b>	<pre> then {   discard;   accept;   count;   log-request; } </pre>
<b>Hierarchy Level</b>	[edit services hcm url-rule url-rule-name term term-num]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Define the HCM term actions.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	<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 384](#)



## then

```
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;
        manual {
            direction (inbound | outbound | bidirectional) {
                authentication {
                    algorithm (hmac-sha-256);
                    key (ascii-text key | hexadecimal key);
                }
                auxiliary-spi spi-value;
                encryption {
                    algorithm algorithm;
                    key (ascii-text key | hexadecimal key);
                }
                protocol (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.

**Usage Guidelines** See *Configuring Match Direction for IPsec Rules*.

**Required Privilege** interface—To view this statement in the configuration.  
**Level** interface-control—To add this statement to the configuration.

## 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 55](#)

## then (Services Stateful Firewall)

<b>Syntax</b>	<pre>then {   (accept   discard   reject);   syslog; }</pre>
<b>Hierarchy Level</b>	[edit services stateful-firewall <a href="#">rule rule-name</a> <a href="#">term term-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define the stateful firewall term actions. You can configure the router to accept, discard, or reject the targeted traffic. The other actions are optional.
<b>Options</b>	<p><b>accept</b>—Accept the traffic and send it on to its destination.</p> <p><b>discard</b>—Do not accept traffic or process it further.</p> <p><b>reject</b>—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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Actions in Stateful Firewall Rules on page 362</a></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> </ul>

## threshold (Services IPsec)

---

<b>Syntax</b>	<code>threshold <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then dead-peer-detection]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	Specify the maximum number of unsuccessful dead peer detection (DPD) requests to be sent before the peer is considered unavailable. (The <b>threshold</b> value is used for IKEv1 security associations (SAs) but not for IKEv2 SAs.)
<b>Options</b>	<b>number</b> —Maximum number of unsuccessful DPD requests to be sent. <b>Range:</b> 1 through 10 <b>Default:</b> 3
<b>Required Privilege Level</b>	security—To view this statement in the configuration. security-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IPsec Rules on page 452</a></li></ul>

## threshold (Services Logging and SYN-Cookie Defenses)

---

<b>Syntax</b>	<code>threshold <i>rate</i>;</code>
<b>Hierarchy Level</b>	[edit services ids <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then logging</a> ], [edit services ids <a href="#">rule</a> <i>rule-name</i> <a href="#">term</a> <i>term-name</i> <a href="#">then syn-cookie</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the threshold for logging or applying SYN-cookie defenses.
<b>Options</b>	<b>rate</b> —Logging threshold number of events per second. <b>rate</b> —SYN-cookie defense number of SYN attacks per second.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Actions in IDS Rules on page 387</a></li></ul>

## traceoptions (Security PKI)

<b>Syntax</b>	<pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i>&gt; &lt;size <i>maximum-file-size</i>&gt;     &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i>; } </pre>
<b>Hierarchy Level</b>	[edit security pki]
<b>Description</b>	Configure security public key infrastructure (PKI) trace options. To specify more than one trace option, include multiple <b>flag</b> statements. Trace option output is recorded in the <code>/var/log/pkid</code> file.
<b>Options</b>	<p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. To include the <b>file</b> statement, you must specify a filename.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file (for example, <b>pkid</b>) reaches its maximum size, it is renamed <b>pkid.0</b>, then <b>pkid.1</b>, 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 <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p><b>flag</b>—Trace operation to perform. To specify more than one trace operation, include multiple flag statements:</p> <ul style="list-style-type: none"> <li><b>all</b>—Trace with all flags enabled.</li> <li><b>certificate-verification</b>—Trace PKI certificate verification events.</li> <li><b>online-crl-check</b>—Trace PKI online certificate revocation list (CRL) events.</li> <li><b>enrollment</b>—PKI certificate enrollment tracing.</li> </ul> <p><b>match <i>regular-expression</i></b>—(Optional) Refine the output to include lines that contain the regular expression.</p> <p><b>size <i>maximum-file-size</i></b>—(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 <b>files <i>number</i></b> option.</p> <p><b>Default:</b> 1024 KB</p> <p><b>world-readable   no-world-readable</b>—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The <b>world-readable</b> option enables any user to read the file. To explicitly set the default behavior, use the <b>no-world-readable</b> option.</p>

<b>Required Privilege</b>	trace—To view this statement in the configuration.
<b>Level</b>	trace-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tracing Junos VPN Site Secure Operations on page 466</a></li></ul>

## traceoptions (Services IPsec VPN)

<b>Syntax</b>	<pre> traceoptions {     file &lt;filename&gt; &lt;files number&gt; &lt;match regular-expression&gt; &lt;size bytes&gt; &lt;world-readable           no-world-readable&gt;;     flag flag;     level level;     no-remote-trace; } </pre>
<b>Hierarchy Level</b>	[edit services ipsec-vpn]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. level option added in Junos OS Release 10.0.
<b>Description</b>	Configure IPsec tracing operations. By default, messages are written to <code>/var/log/kmd</code> .
<b>Options</b>	<p><b>files <i>number</i></b>—Maximum number of trace data files.  <b>Range:</b> 2 through 1000</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform:</p> <ul style="list-style-type: none"> <li>• <b>all</b>—Trace everything.</li> <li>• <b>certificates</b>—Trace certificates that apply to the IPsec service set.</li> <li>• <b>database</b>—Trace security associations database events.</li> <li>• <b>general</b>—Trace general events.</li> <li>• <b>ike</b>—Trace IKE module processing.</li> <li>• <b>parse</b>—Trace configuration processing.</li> <li>• <b>policy-manager</b>—Trace policy manager processing.</li> <li>• <b>routing-socket</b>—Trace routing socket messages.</li> <li>• <b>snmp</b>—Trace SNMP operations.</li> <li>• <b>timer</b>—Trace internal timer events.</li> </ul> <p><b>level <i>level</i></b>—Key management process (kmd) tracing level. The following values are supported:</p> <ul style="list-style-type: none"> <li>• <b>all</b>—Match all levels.</li> <li>• <b>error</b>—Match error conditions.</li> <li>• <b>info</b>—Match informational messages.</li> <li>• <b>notice</b>—Match conditions that should be handled specially.</li> <li>• <b>verbose</b>—Match verbose messages.</li> <li>• <b>warning</b>—Match warning messages.</li> </ul>

**size bytes**—Maximum trace file size.

<b>Required Privilege</b>	interface—To view this statement in the configuration.
<b>Level</b>	interface-control—To add this statement to the configuration.



## traceoptions (Services L2TP)

<b>Syntax</b>	<pre> traceoptions {   debug-level <i>level</i>;   file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i> &gt; &lt;size <i>maximum-file-size</i>&gt;     &lt;world-readable   no-world-readable&gt;;   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>
<b>Hierarchy Level</b>	[edit services l2tp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define tracing operations for L2TP processes.
<b>Options</b>	<p><b>debug-level <i>level</i></b>—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"> <li><b>detail</b>—Trace detailed debug information.</li> <li><b>error</b>—Trace error information.</li> <li><b>packet-dump</b>—Trace packet decoding information.</li> </ul> <p><b>file <i>filename</i></b>—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 <code>/var/log</code>.</p> <p><b>files <i>number</i></b>—(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 <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>filter</b>—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"> <li><b>protocol <i>name</i></b>—One of the following protocols; this option does not apply to L2TP on MX Series routers: <ul style="list-style-type: none"> <li><b>l2tp</b></li> </ul> </li> </ul>

- **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.

<b>Required Privilege Level</b>	trace—To view this statement in the configuration.
	trace-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tracing L2TP Operations on page 689</a></li></ul>
	<ul style="list-style-type: none"><li>• <i>Tracing L2TP Operations for Subscriber Access</i></li></ul>

## traceoptions (Services Logging)

<b>Syntax</b>	<pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable       no-world-readable&gt;;     flag <i>flag</i>;     no-remote-trace; } </pre>
<b>Hierarchy Level</b>	<p>[edit services <a href="#">adaptive-services-pics</a>],  [edit services <a href="#">logging</a>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.  <b>file</b> option added in Release 8.0.</p>
<b>Description</b>	<p>Configure Adaptive Services or Multiservices PIC tracing operations. The messages are output to <b>/var/log/serviced</b>.</p>
<b>Options</b>	<p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. All files are placed in the directory <b>/var/log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, 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 <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files  <b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform:</p> <ul style="list-style-type: none"> <li>• <b>all</b>—Trace everything.</li> <li>• <b>command-queued</b>—Trace command enqueue events.</li> <li>• <b>config</b>—Trace configuration events.</li> <li>• <b>handshake</b>—Trace handshake events.</li> <li>• <b>init</b>—Trace initialization events.</li> <li>• <b>interfaces</b>—Trace interface events.</li> <li>• <b>mib</b>—Trace GGSN SNMP MIB events.</li> <li>• <b>removed-client</b>—Trace client cleanup events.</li> <li>• <b>show</b>—Trace CLI command servicing.</li> </ul> <p><b>match <i>regex</i></b>—(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.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration.
	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tracing Services PIC Operations on page 27</a></li></ul>

## translated

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">rule</a> rule-name <a href="#">term</a> term-name then]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define properties for translated traffic.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li> </ul>

## transport

<b>Syntax</b>	transport [ <i>transport-protocols</i> ];
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">pool</a> nat-pool-name <a href="#">pgcp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2.
<b>Description</b>	Configure the BGF to select a NAT pool based on transport protocol type.
<b>Options</b>	<p>[ <i>transport-protocol</i> ]—One or more transport protocols.</p> <p><b>Values:</b> rtp-avp, tcp, udp</p> <p><b>Syntax:</b> One or more protocols. If you specify more than one protocol, you must enclose all protocols in brackets.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## trigger-link-failure

---

<b>Syntax</b>	<code>trigger-link-failure <i>interface-name</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <code>lsq-fpc/pic/port</code> <a href="#">lsq-failure-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4.
<b>Description</b>	List of SONET interfaces connected to the LSQ interface that can implement Automatic Protection Switching (APS) if the Link Services IQ PIC fails.
<b>Options</b>	<i>interface-name</i> —Name of SONET interface.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Association between LSQ and SONET Interfaces on page 590</a></li></ul>

## translated-port

---

<b>Syntax</b>	<code>translated-port <i>port id</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">services</a> nat <a href="#">port-forwarding</a> <i>map-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	Specify the port to which all traffic will be translated.
<b>Options</b>	<i>port id</i> —The port number to which traffic will be translated.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">port-forwarding on page 852</a></li><li>• <a href="#">destined-port on page 754</a></li></ul>



## translation-type

<b>Syntax</b>	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)
<b>Hierarchy Level</b>	[edit <b>services</b> nat <b>rule</b> <i>rule-name</i> <b>term</b> <i>term-name</i> <b>then</b> <b>translated</b> ]
<b>Release Information</b>	<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"> <li>• <b>basic-nat44</b></li> <li>• <b>basic-nat66</b></li> <li>• <b>basic-nat-pt</b></li> <li>• <b>deterministic-napt44</b></li> <li>• <b>dnat-44</b></li> <li>• <b>dynamic-nat44</b></li> <li>• <b>napt-44</b></li> <li>• <b>napt-66</b></li> <li>• <b>napt-pt</b></li> <li>• <b>stateful-nat64</b></li> </ul> <p><b>twice-basic-nat-44</b> option introduced in Junos OS Release 11.4.</p> <p><b>twice-dynamic-nat-44</b> option introduced in Junos OS Release 11.4.</p> <p><b>twice-napt-44</b> option introduced in Junos OS Release 11.4.</p> <p><b>deterministic-napt44</b> option introduced in Junos OS Release 12.1.</p> <p><b>nptv6</b> option introduced in Junos OS Release 15.1</p>
<b>Description</b>	Specify the NAT translation types.
<b>Options</b>	<ul style="list-style-type: none"> <li>• <b>basic-nat44</b>—Translate the source address statically (IPv4 to IPv4).</li> <li>• <b>basic-nat66</b>—Translate the source address statically (IPv6 to IPv6).</li> <li>• <b>basic-nat-pt</b>—Translate the addresses of IPv6 hosts as they originate sessions to the IPv4 hosts in the external domain. The <b>basic-nat-pt</b> option is always implemented with DNS ALG.</li> <li>• <b>deterministic-napt44</b>—Translate as <b>napt-44</b>, and use deterministic port block allocation for port translation.</li> <li>• <b>dnat-44</b>—Translate the destination address statically (IPv4 to IPv4).</li> <li>• <b>dynamic-nat44</b>—Translate only the source address by dynamically choosing the NAT address from the source address pool.</li> </ul>

- **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).



**NOTE:** Starting with Junos OS Release 15.1, the twice NAT functionality (**twice-basic-nat-44**, **twice-dynamic-nat-44**, and **twice-dynamic-napt-44** options) is supported on MX Series routers with MS-MPCs and MS-MICs.

---

- **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.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration.
	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Network Address Translation Rules Overview on page 55</a></li></ul>

## trusted-ca

---

<b>Syntax</b>	<code>trusted-ca <i>ca-profile-name</i>;</code>
<b>Hierarchy Level</b>	[edit services <a href="#">service-set</a> <i>service-set-name</i> <a href="#">ipsec-vpn-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Identify one or more trusted IPsec certification authorities.
<b>Options</b>	<i>ca-profile-name</i> —Name of certification authority profile, which is configured at the [edit <a href="#">security pki</a> ] hierarchy level.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPsec Service Sets on page 460</a></li> </ul>

## ttl-threshold

---

<b>Syntax</b>	<code>ttl-threshold <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">applications application</a> <i>application-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the traceroute time-to-live (TTL) threshold value. This value sets the acceptable level of network penetration for trace routing.
<b>Options</b>	<i>number</i> —TTL threshold value.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">ALG Descriptions on page 299</a></li> <li>• <a href="#">Configuring the TTL Threshold on page 342.</a></li> <li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li> <li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li> </ul>

## 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 679](#)
  - [Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces](#)

## tunnel-mtu (Services IPsec VPN)

<b>Syntax</b>	tunnel-mtu <i>bytes</i> ;
<b>Hierarchy Level</b>	[edit services ipsec-vpn <a href="#">rule rule-name</a> <a href="#">term term-name</a> <a href="#">then</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Maximum transmission unit (MTU) size for IPsec tunnels. This defines the maximum size of an IP packet, including the IPsec overhead.
<b>Options</b>	<p><i>bytes</i>—MTU size.</p> <p><b>Default:</b> 1500 bytes</p> <p><b>Range:</b> 256 through 9192 bytes</p>
<div>  <p><b>NOTE:</b> Clear the IPsec SA in tunnel-mtu to accommodate Jumbo frames larger than 1500 bytes.</p> </div>	
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

## tunnel-mtu (Services Service Set)

---

<b>Syntax</b>	<code>tunnel-mtu bytes;</code>
<b>Hierarchy Level</b>	<code>[edit services service-set <i>service-set-name</i> ipsec-vpn-options]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0.
<b>Description</b>	<p>Maximum transmission unit (MTU) size for IPsec tunnels. This statement is useful for dynamic endpoint tunnels for which you cannot configure the <b>tunnel-mtu</b> statement at the <code>[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then]</code> 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 <b>tunnel-mtu</b> statement at the <code>[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then]</code> hierarchy level.</p>
	<div> <b>NOTE:</b> The <b>tunnel-mtu</b> setting at the <code>[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then]</code> hierarchy level overrides the value specified at the <code>[edit services service-set <i>service-set-name</i> ipsec-vpn-options]</code> hierarchy level.</div>
<b>Options</b>	<p><b>bytes</b>—MTU size.</p> <p><b>Default:</b> 1500 bytes</p> <p><b>Range:</b> 256 through 9192 bytes</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>mtu</i></li><li>• <a href="#">Configuring IPsec Service Sets on page 460</a></li><li>• <a href="#">Configuring IPsec Rules on page 452</a></li></ul>

## tunnel-timeout

---

<b>Syntax</b>	tunnel-timeout <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit services l2tp <b>tunnel-group</b> <i>name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the maximum downtime for an L2TP tunnel, after which the tunnel is terminated because the connection is presumed to have been lost.
<b>Options</b>	<i>seconds</i> —Interval after which the tunnel is terminated if no data can be sent. <b>Default:</b> 120 seconds
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Timers for L2TP Tunnels on page 681</a></li> <li>• <a href="#">Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces</a></li> </ul>

## url

---

<b>Syntax</b>	url <i>url_identifier</i> ;
<b>Hierarchy Level</b>	[edit services hcm url-rule <i>url-rule-name</i> term <i>term-num</i> from]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Specify an integer that uniquely identifies a particular URL definition within a term.
<b>Options</b>	<i>url_identifier</i> —URL identifier number. <b>Range:</b> 1 through 32,767 <b>Default:</b> If no value is added, the default value is 1.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## url-list

---

<b>Syntax</b>	<code>url-list <i>url-list-name</i> ;</code>
<b>Hierarchy Level</b>	<code>[edit services hcm url-rule <i>url-rule-name</i> term <i>term-num</i> from]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	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.
<b>Options</b>	<i>url-list-name</i> —Name of the previously defined URL list.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## url-rule

---

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	<code>[edit services hcm]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Specify the name of the URL rule.
<b>Options</b>	<i>url-rule-name</i> —Name of the URL rule.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.



## url-rule-set

<b>Syntax</b>	<code>url-rule-set <i>url-rule-set-name</i> {     url-rule <i>rule1</i>;     url-rule <i>rule2</i>; }</code>
<b>Hierarchy Level</b>	<code>[edit services hcm url-rule <i>url-rule-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2.
<b>Description</b>	Specify the name of the rule set. A rule set is a collection of rules ordered in the sequence in which they are entered.
<b>Options</b>	<i>url-rule-set-name</i> —Name of the collection of URL rules that constitute this rule set.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## unit (Aggregated Multiservices)

<b>Syntax</b>	<code>unit <i>interface-unit-number</i> {     family <i>family</i>; }</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	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.
<b>Options</b>	<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

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Aggregated Multiservices Interfaces on page 643</a></li> <li>• <a href="#">Example: Configuring an Aggregated Multiservices Interface (AMS) on page 652</a></li> <li>• <a href="#">interfaces on page 801</a></li> </ul>

## unit (Interfaces)

---

<b>Syntax</b>	<pre>unit <i>logical-unit-number</i> {     family inet {         address <i>address</i> {         }         service {             input {                 [ <i>service-set</i> <i>service-set-name</i> &lt;<i>service-filter</i> <i>filter-name</i>&gt; ];                 <i>post-service-filter</i> <i>filter-name</i>;             }             output {                 [ <i>service-set</i> <i>service-set-name</i> &lt;<i>service-filter</i> <i>filter-name</i>&gt; ];             }         }         <i>service-domain</i> (inside   outside);     } }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.
<b>Options</b>	<p><i>logical-unit-number</i>—Number of the logical unit.</p> <p><b>Range:</b> 0 through 16,384</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces.</li></ul>

## 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 666](#)

## uuid

---

<b>Syntax</b>	<code>uuid <i>hex-value</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">applications application application-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the Universal Unique Identifier (UUID) for DCE RPC objects.
<b>Options</b>	<i>hex-value</i> —Hexadecimal value.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">ALG Descriptions on page 299</a></li><li>• <a href="#">Configuring a Universal Unique Identifier on page 342</a></li><li>• <a href="#">Examples: Configuring Application Protocols on page 343</a></li><li>• <a href="#">Verifying the Output of ALG Sessions on page 344</a></li></ul>

## v6rd

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit services softwire <a href="#">softwire-concentrator</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	Configure settings for a 6rd concentrator used to process IPv6 packets encapsulated in IPv4 packets.
<b>Options</b>	<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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a 6rd Softwire Concentrator on page 255</a></li> </ul>

## version (IKE)

---

<b>Syntax</b>	<code>version ( 1   2 );</code>
<b>Hierarchy Level</b>	<code>[edit services ipsec-vpn ike policy <i>policy-name</i>],</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	Configure the Internet Key Exchange (IKE) version that is used to negotiate dynamic SAs for IPsec.
<b>Options</b>	1—Uses IKEv1. 2—Uses IKEv2.



**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.

<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IKE Policies on page 439</a></li></ul>

## video

---

<b>Syntax</b>	<pre>video {     dscp (<i>alias</i>   <i>bits</i>);     forwarding-class (Services CoS) <i>class-name</i>; }</pre>
<b>Hierarchy Level</b>	<code>[edit services (CoS) cos application-profile <i>profile-name</i> sip]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3.
<b>Description</b>	Set the appropriate <b>dscp</b> and <b>forwarding-class</b> values for SIP video traffic.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Application Profiles for Use as CoS Rule Actions on page 559</a></li></ul>

## video (Application Profile)

---

<b>Syntax</b>	<pre>video {     dscp (alias   bits);     forwarding-class class-name; }</pre>
<b>Hierarchy Level</b>	[edit services cos <a href="#">application-profile</a> <i>profile-name</i> sip]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3.
<b>Description</b>	Set the appropriate <b>dscp</b> and <b>forwarding-class</b> values for SIP video traffic.
<b>Default</b>	By default, the system will not alter the DSCP or forwarding class for SIP video traffic.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Application Profiles</i></li> <li>• <a href="#">voice (Application Profile) on page 958</a></li> </ul>

## voice

---

<b>Syntax</b>	<pre>voice {     dscp (alias   bits);     forwarding-class (Services CoS) class-name; }</pre>
<b>Hierarchy Level</b>	[edit services (CoS) cos <a href="#">application-profile</a> <i>profile-name</i> sip]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3.
<b>Description</b>	Set the appropriate <b>dscp</b> and <b>forwarding-class</b> values for SIP voice traffic.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Application Profiles for Use as CoS Rule Actions on page 559</a></li> </ul>

## voice (Application Profile)

---

<b>Syntax</b>	<pre>voice {     dscp (<i>alias</i>   <i>bits</i>);     forwarding-class <i>class-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit services cos <a href="#">application-profile</a> <i>profile-name</i> sip]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3.
<b>Description</b>	Set the appropriate <b>dscp</b> and <b>forwarding-class</b> values for SIP voice traffic.
<b>Default</b>	By default, the system will not alter the DSCP or forwarding class for SIP voice traffic.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Application Profiles</i></li><li>• <a href="#">video (Application Profile) on page 957</a></li></ul>

## warm-standby

---

<b>Syntax</b>	<pre>warm-standby;</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>rls</i> <i>number</i> <a href="#">redundancy-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.0.
<b>Description</b>	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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 592</a></li></ul>



## CHAPTER 44

# 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 inline software 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 sessions
- 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 hcm statistics
- show services ids

- `show services inline nat pool`
- `show services inline nat statistics`
- `show services inline software 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 deterministic-nat internal-host`
- `show services nat deterministic-nat nat-port-block`
- `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 integrity-drops`
- `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 sessions`
- `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

---

<b>Syntax</b>	clear services cos statistics <interface <i>interface-name</i> > <service-set <i>service-set-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 8.1.
<b>Description</b>	Clear statistics for class-of-service (CoS) code point bit patterns and forwarding classes as configured in CoS services for the AS PIC.
<b>Options</b>	<p><b>none</b>—Clear all services CoS statistics.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear statistics for the specified interface only.</p> <p><b>service-set <i>service-set-name</i></b>—(Optional) Clear statistics for the specified service set only.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">clear services cos statistics on page 963</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	clear services crtp statistics <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Clear Compressed Real-Time Transport Protocol (CRTP) flow statistics.
<b>Options</b>	<b>none</b> —Clear CRTP flow statistics on all interfaces.  <b>interface <i>interface-name</i></b> —(Optional) Clear CRTP flow statistics for the specified interface. On M Series and T Series routers, a link services IQ ( <b>lsq-<i>fpc/pic/port</i></b> ) or redundant link services IQ ( <b>rlsq-<i>fpc/pic/port</i></b> ) interface.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">clear services crtp statistics on page 964</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	clear services ids <interface <i>interface-name</i> > <service-set <i>service-set-name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Clear intrusion detection service (IDS) events.
<b>Options</b>	<p><b>none</b>—Clear all IDS events for all adaptive services interfaces for all service sets, and clear and reset IDS.</p> <p><b>interface <i>interface-name</i></b>—(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><b>service-set <i>service-set-name</i></b>—(Optional) Clear all IDS events for a particular service set.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">clear services ids on page 965</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	<code>clear services ids destination-table</code> <code>&lt;destination-prefix <i>destination-prefix-name</i>&gt;</code> <code>&lt;interface <i>interface-name</i>&gt;</code> <code>&lt;service-set <i>service-set-name</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Clear the intrusion detection service (IDS) events for a particular address that might be under attack.
<b>Options</b>	<p><b>none</b>—Clear the attack destination address table.</p> <p><b>destination-prefix <i>destination-prefix-name</i></b>—(Optional) Clear the attack destination table for a particular destination prefix.</p> <p><b>interface <i>interface-name</i></b>—(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><b>service-set <i>service-set-name</i></b>—(Optional) Clear the attack destination table for a particular service set.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">clear services ids destination-table on page 966</a>
<b>Output Fields</b>	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

<b>Syntax</b>	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> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Clear the intrusion detection service (IDS) attack source and destination address pair table.
<b>Options</b>	<p><b>none</b>—Clear the attack source and destination address pair table.</p> <p><b>destination-prefix <i>destination-prefix-name</i></b>—(Optional) Clear the attack source and destination address pair table for a particular destination prefix.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear the attack destination table for a particular interface. On M Series and T Series routers, the <i>interface-name</i> can be <b>sp-fpc/pic/port</b> or <b>rspnumber</b>.</p> <p><b>service-set <i>service-set-name</i></b>—(Optional) Clear the attack source and destination address pair table for a particular service set.</p> <p><b>source-prefix <i>source-prefix-name</i></b>—(Optional) Clear the attack source and destination address pair table for a particular source prefix.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">clear services ids pair-table on page 967</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	<code>clear services ids source-table</code> <code>&lt;interface <i>interface-name</i>&gt;</code> <code>&lt;service-set <i>service-set-name</i>&gt;</code> <code>&lt;source-prefix <i>source-prefix-name</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Clear all intrusion detection service (IDS) events for addresses that are suspected attackers.
<b>Options</b>	<p><b>none</b>—Clear the attack source address table.</p> <p><b>interface <i>interface-name</i></b>—(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><b>service-set <i>service-set-name</i></b>—(Optional) Clear the attack source address table for a particular service set.</p> <p><b>source-prefix <i>source-prefix-name</i></b>—(Optional) Clear the attack source address table for a particular source prefix.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">clear services ids source-table on page 968</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	clear services inline nat pool <i>pool-name</i>
<b>Release Information</b>	Command introduced in Junos OS Release 11.4.
<b>Description</b>	Clear global inline NAT statistics.
<b>Options</b>	<b>pool-name</b> —Name of the NAT pool for which statistic are cleared.
<b>Required Privilege Level</b>	clear
<b>List of Sample Output</b>	<a href="#">clear services inline nat pool on page 969</a>
<b>Output Fields</b>	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

---


<b>Syntax</b>	clear services inline nat statistics <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 11.4.
<b>Description</b>	Clear global inline NAT statistics.
<b>Options</b>	<b>interface <i>interface-name</i></b> —(Optional) Clear inline NAT statistics for the specified interface only.
<b>Required Privilege Level</b>	clear
<b>List of Sample Output</b>	<a href="#">clear services inline nat statistics on page 970</a>
<b>Output Fields</b>	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 inline software statistics

<b>Syntax</b>	clear services inline software statistics <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 13.3R3.
<b>Description</b>	Clear global inline software statistics.
<div>  <b>NOTE:</b> The following two limitations apply to the clearing of data plane statistics using the clear services inline software statistics command:           <ul style="list-style-type: none"> <li>• When traffic is continuously flowing and the counters are being updated in the data plane, none of the statistical values except the counter for 6rd decapsulation errors is reset.</li> <li>• When you delete the software concentrator or the service set associated with an inline services (si-) interface, the counter for 6rd decapsulation errors might display all the previously accumulated values.</li> </ul> </div>	
<b>Options</b>	<b>interface <i>interface-name</i></b> —(Optional) Clear inline software statistics for the specified interface only.
<b>Required Privilege Level</b>	clear
<b>List of Sample Output</b>	<a href="#">clear services inline software statistics on page 971</a>
<b>Output Fields</b>	When you enter this command, the global inline software statistics are cleared. There is no specific output.

### Sample Output

#### clear services inline software statistics

```
user@host> clear services inline software statistics
```

## clear services ipsec-vpn certificates

---

<b>Syntax</b>	clear services ipsec-vpn certificates (all   service-set <i>service-set</i> ) <certificate-cache-entry <i>number</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	(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.
<b>Options</b>	<b>all</b> —Delete digital certificates for all service sets.  <b>service-set <i>service-set</i></b> —Delete digital certificates for the specified service set.
<b>Required Privilege Level</b>	clear
<b>List of Sample Output</b>	<a href="#">clear services ipsec-vpn certificates all on page 972</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	clear services ipsec-vpn ike security-associations <peer-address-name> <service-set service-set-name>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. service-set option added in Junos OS Release 8.5.
<b>Description</b>	(Adaptive services interfaces only) Clear Internet Key Exchange (IKE) security associations.
<b>Options</b>	<p><b>peer-address-name</b>—(Optional) Clear only the security association specified by the peer address.</p> <p><b>service-set service-set-name</b>—(Optional) Clear only the security association specified by the service-set name.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show services ipsec-vpn ike security-associations on page 1120</a></li> </ul>
<b>Output Fields</b>	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

---

<b>Syntax</b>	<code>clear services ipsec-vpn security-associations</code> <code>&lt;peer-address-name&gt;</code> <code>&lt;remote-gateway remote-gateway-address&gt;</code> <code>&lt;service-set-name&gt;</code> <code>&lt;tunnel-index tunnel-index-number&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. <b>remote-gateway</b> , <b>service-set-name</b> , and <b>tunnel-index</b> options added in Junos OS Release 8.4.
<b>Description</b>	(Adaptive services interfaces only) Clear IP Security (IPsec) security associations. You can combine the options for greater specificity.
<b>Options</b>	<p><b>peer-address-name</b>—(Optional) Clear only the security association specified by the peer address.</p> <p><b>remote-gateway remote-gateway-address</b>—(Optional) Clear only the security association specified by the remote gateway address.</p> <p><b>service-set-name</b>—(Optional) Clear only the security association specified by the service-set name.</p> <p><b>tunnel-index tunnel-index-number</b>—(Optional) Clear only the security association specified by the tunnel index number.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show services ipsec-vpn ipsec security-associations on page 1124</a></li></ul>
<b>Output Fields</b>	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

---

<b>Syntax</b>	clear services ipsec-vpn ipsec statistics <remote-gateway <i>address</i> > <service-set <i>service-set-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 8.1.
<b>Description</b>	(Adaptive services interface only) Clear IP Security (IPsec) statistics.
<b>Options</b>	<b>remote-gateway <i>address</i></b> —(Optional) Clear statistics for the specified remote system.  <b>service-set <i>service-set-name</i></b> —(Optional) Clear statistics for the specified service set.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show services ipsec-vpn ipsec statistics on page 1128</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear services ipsec-vpn ipsec statistics on page 975</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	<code>clear services l2tp destination</code> <code>&lt;all   local-gateway <i>gateway-address</i>   peer-gateway <i>gateway-address</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.4.
<b>Description</b>	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.
<b>Options</b>	<p><b>all</b>—Close all L2TP destinations.</p> <p><b>local-gateway <i>gateway-address</i></b>—Clear only the L2TP destinations and all tunnels and sessions associated with the specified local gateway address.</p> <p><b>peer-gateway <i>gateway-address</i></b>—Clear only the L2TP destinations and all tunnels and sessions associated with the peer gateway with the specified address.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show services l2tp destination</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear services l2tp destination all on page 976</a>
<b>Output Fields</b>	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

<b>Syntax</b>	clear services l2tp destination statistics <all   local-gateway <i>gateway-address</i>   peer-gateway <i>gateway-address</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 13.1.
<b>Description</b>	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.
<b>Options</b>	<p><b>all</b>—Clear all statistics associated with the L2TP destinations.</p> <p><b>local-gateway <i>gateway-address</i></b>—Clear statistics related to L2TP destination and all tunnels and sessions associated with the specified local gateway address.</p> <p><b>peer-gateway <i>gateway-address</i></b>—Clear statistics related to L2TP destination and all tunnels and sessions associated with the specified peer gateway address.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>show services l2tp destination</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear services l2tp destination statistics on page 977</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	clear services l2tp multilink (all <statistics>   bundle-id <i>number</i> <statistics>   statistics (all   bundle-id <i>number</i> ))
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(M10i and M7i routers only) Close Layer 2 Tunneling Protocol (L2TP) multilink sessions or clear session statistics.
<b>Options</b>	<p><b>all &lt;statistics&gt;</b>—Close all L2TP multilink sessions or clear statistics for all L2TP multilink sessions.</p> <p><b>bundle-id <i>number</i> &lt;statistics&gt;</b>—L2TP multilink bundle ID. The value is an internally generated number from 1 to 65535. Close the specified L2TP multilink session, or using the <b>statistics</b> keyword with this option, clear statistics for the specified session.</p> <p><b>statistics (all   bundle-id <i>number</i>)</b>—Clear all session statistics or clear statistics for the specified multilink bundle ID.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">L2TP Services Configuration Overview on page 676</a></li><li>• <a href="#">L2TP Minimum Configuration on page 677</a></li><li>• <a href="#">show services l2tp multilink on page 1135</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear services l2tp multilink statistics all on page 978</a>
<b>Output Fields</b>	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

<b>Syntax</b>	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> )
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(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.
<b>Options</b>	<p><b>all</b>—Close all L2TP sessions.</p> <p><b>interface <i>interface-name</i></b>—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"> <li>• <b>si-<i>fpc/pic/port</i></b>—MPCs on MX Series routers only. This option is not available for L2TP on M Series routers.</li> <li>• <b>sp-<i>fpc/pic/port</i></b>—AS or Multiservices PICs on M7i, M10i, and M120 routers only. This option is not available for L2TP on MX Series routers.</li> </ul> <p><b>local-gateway <i>gateway-address</i></b>—Clear only the L2TP sessions associated with the specified local gateway address.</p> <p><b>local-gateway-name <i>gateway-name</i></b>—Clear only the L2TP sessions associated with the specified local gateway name.</p> <p><b>local-session-id <i>session-id</i></b>—Clear only the L2TP sessions with this identifier for the local endpoint of the L2TP session.</p> <p><b>local-tunnel-id <i>tunnel-id</i></b>—Clear only the L2TP sessions associated with the specified local tunnel identifier.</p> <p><b>peer-gateway <i>gateway-address</i></b>—Clear only the L2TP sessions associated with the peer gateway with the specified address.</p> <p><b>peer-gateway-name <i>gateway-name</i></b>—Clear only the L2TP sessions associated with the peer gateway with the specified name.</p> <p><b>tunnel-group <i>group-name</i></b>—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><b>user <i>username</i></b>—(M Series routers only) Clear only the L2TP sessions for the specified username.</p>
<b>Required Privilege Level</b>	clear

- Related Documentation**
- [L2TP Services Configuration Overview on page 676](#)
  - [L2TP Minimum Configuration on page 677](#)
  - [clear services l2tp session statistics on page 981](#)
  - [show services l2tp session on page 1143](#)

**List of Sample Output**    [clear services l2tp session on page 980](#)  
                                  [clear services l2tp session interface on page 980](#)

**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

<b>Syntax</b>	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> )
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Support for MX Series routers added in Junos OS Release 10.4.
<b>Description</b>	(M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Clear statistics for Layer 2 Tunneling Protocol (L2TP) sessions.
<b>Options</b>	<p><b>all</b>—Clear statistics for all L2TP sessions.</p> <p><b>interface <i>interface-name</i></b>—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"> <li>• <b>si-<i>fpc/pic/port</i></b>—MPCs on MX Series routers only. This option is not available for L2TP on M Series routers.</li> <li>• <b>sp-<i>fpc/pic/port</i></b>—AS or Multiservices PICs on M7i, M10i, and M120 routers only. This option is not available for L2TP on MX Series routers.</li> </ul> <p><b>local-gateway <i>gateway-address</i></b>—Clear statistics for only the L2TP sessions associated with the local gateway with the specified address.</p> <p><b>local-gateway-name <i>gateway-name</i></b>—Clear statistics for only the L2TP sessions associated with the local gateway with the specified name.</p> <p><b>local-session-id <i>session-id</i></b>—Clear statistics for only the L2TP sessions with this identifier for the local endpoint of the L2TP session.</p> <p><b>local-tunnel-id <i>tunnel-id</i></b>—Clear statistics for only the L2TP sessions associated with the specified local tunnel identifier.</p> <p><b>peer-gateway <i>gateway-address</i></b>—Clear statistics for only the L2TP sessions associated with the peer gateway with the specified address.</p> <p><b>peer-gateway-name <i>gateway-name</i></b>—Clear statistics for only the L2TP sessions associated with the peer gateway with the specified name.</p> <p><b>tunnel-group <i>group-name</i></b>—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><b>user <i>username</i></b>—Clear statistics for only the L2TP sessions for the specified username. This option is not available for L2TP LAC on MX Series routers.</p>
<b>Required Privilege Level</b>	view

- Related Documentation**
- [L2TP Services Configuration Overview on page 676](#)
  - [L2TP Minimum Configuration on page 677](#)
  - [clear services l2tp session on page 979](#)
  - [show services l2tp session on page 1143](#)

**List of Sample Output**   [clear services l2tp session statistics all on page 982](#)

**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

<b>Syntax</b>	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> )
<b>Release Information</b>	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.
<b>Description</b>	(M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Clear Layer 2 Tunneling Protocol (L2TP) tunnels.
<b>Options</b>	<p><b>all</b>—Clear all L2TP tunnels.</p> <p><b>sp-fpc/pic/port</b>—(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><b>local-gateway gateway-address</b>—Clear only the L2TP tunnels associated with the local gateway with the specified address.</p> <p><b>local-gateway-name gateway-name</b>—Clear only the L2TP tunnels associated with the local gateway with the specified name.</p> <p><b>local-tunnel-id tunnel-id</b>—Clear only the L2TP tunnels that have the specified local tunnel identifier.</p> <p><b>peer-gateway gateway-address</b>—Clear only the L2TP tunnels associated with the peer gateway with the specified address.</p> <p><b>peer-gateway-name gateway-name</b>—Clear only the L2TP tunnels associated with the peer gateway with the specified name.</p> <p><b>tunnel-group group-name</b>—Clear only the L2TP tunnels in the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">L2TP Services Configuration Overview on page 676</a></li> <li>• <a href="#">L2TP Minimum Configuration on page 677</a></li> <li>• <a href="#">clear services l2tp tunnel statistics on page 985</a></li> <li>• <a href="#">show services l2tp tunnel on page 1156</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear services l2tp tunnel on page 984</a>
<b>Output Fields</b>	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

<b>Syntax</b>	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> )
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Support for MX Series routers added in Junos OS Release 10.4.
<b>Description</b>	(M10i and M7i routers: LNS only. MX Series routers: LAC only.) Clear statistics for Layer 2 Tunneling Protocol (L2TP) tunnels.
<b>Options</b>	<p><b>all</b>—Clear statistics for all L2TP tunnels.</p> <p><b>interface <i>sp-fpc/pic/port</i></b>—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><b>local-gateway <i>gateway-address</i></b>—Clear statistics for only the L2TP tunnels associated with the local gateway with the specified address.</p> <p><b>local-gateway-name <i>gateway-name</i></b>—Clear statistics for only the L2TP tunnels associated with the local gateway with the specified name.</p> <p><b>local-tunnel-id <i>tunnel-id</i></b>—Clear statistics for only the L2TP tunnels that have the specified local tunnel identifier.</p> <p><b>peer-gateway <i>gateway-address</i></b>—Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified address.</p> <p><b>peer-gateway-name <i>gateway-name</i></b>—Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified name.</p> <p><b>tunnel-group <i>group-name</i></b>—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>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">L2TP Services Configuration Overview on page 676</a></li> <li>• <a href="#">L2TP Minimum Configuration on page 677</a></li> <li>• <a href="#">clear services l2tp tunnel on page 983</a></li> <li>• <a href="#">show services l2tp tunnel on page 1156</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear services l2tp tunnel statistics all on page 986</a>
<b>Output Fields</b>	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

<b>Syntax</b>	clear services nat flows <b4address <b><i>b4address</i></b> > <service-set <i>service-set</i> > <subscriber <i>subscriber-address</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 14.1.
<b>Description</b>	Clear NAT flows.
<b>Options</b>	<p><b>none</b>—Clear all NAT flows.</p> <p><b>b4address <i>b4address</i></b>—(Optional) Clear NAT flows for a particular B4 address.</p> <p><b>service-set <i>service-set</i></b>—(Optional) Clear NAT flows for a particular service set.</p> <p><b>subscriber <i>ip</i></b>—(Optional) Clear NAT flows for a particular subscriber, identified by IPv4 address.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	
<b>List of Sample Output</b>	<a href="#">clear services nat flows subscriber (IPv4 address) on page 987</a>
<b>Output Fields</b>	<a href="#">Table 27 on page 987</a> lists the output fields for the <b>clear services nat flows</b> command. Output fields are listed in the approximate order in which they appear.

**Table 27: clear services nat flows Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of a services interface.
<b>Service set</b>	Name of the service set from which flows are being cleared.
<b>Flows removed</b>	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

<b>Syntax</b>	clear services nat mappings <app> <eim> <pcp> <service-set service-set>
<b>Release Information</b>	Command introduced in Junos OS Release 14.1.
<b>Description</b>	Clear NAT mappings.
<b>Options</b>	<p><b>none</b>—Clear all NAT mappings.</p> <p><b>app</b>—(Optional) Clear address-pooling paired NAT mappings.</p> <p><b>eim</b>—(Optional) Clear endpoint-independent NAT mappings.</p> <p><b>pcp</b>—(Optional) Clear Port Control Protocol NAT mappings.</p> <p><b>service-set service-set</b>—(Optional) Clear NAT mappings for a specified service set..</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show services nat mappings on page 1171</a></li> <li>• <a href="#">clear services nat mappings app on page 990</a></li> <li>• <a href="#">clear services nat mappings eim on page 991</a></li> <li>• <a href="#">clear services nat mappings pcp on page 993</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear services nat mappings on page 989</a>
<b>Output Fields</b>	Table 28 on page 988 lists the output fields for the <b>clear services nat mappings</b> command. Output fields are listed in the approximate order in which they appear.

**Table 28: clear services nat mappings Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of a services interface.
<b>Service set</b>	Name of the service set from which flows are being cleared.
<b>Mappings removed</b>	Number of mappings removed.
<b>Flows removed</b>	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

<b>Syntax</b>	clear services nat mappings app <b4address <i>b4address/prefix</i> > <service-set <i>service-set</i> > <subscriber <i>subscriber-ipv4-address</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 14.1.
<b>Description</b>	Clear NAT mappings for address pooling paired (app).
<b>Options</b>	<p><b>none</b>—Clear all NAT app mappings.</p> <p><b>b4address <i>b4address/prefix</i></b>—(Optional) Clear NAT APP mappings for a particular subscriber <i>b4address/prefix</i></p> <p><b>service-set <i>service-set</i></b>—(Optional) Clear NAT APP mappings for a specified service set..</p> <p><b>subscriber <i>subscriber-ipv4-address/prefix</i></b>—(Optional) Clear NAT APP mappings for a particular subscriber <i>ipv4-address/prefix</i></p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show services nat mappings on page 1171</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear services nat mappings app on page 990</a>
<b>Output Fields</b>	Table 29 on page 990 lists the output fields for the <b>clear services nat mappings app</b> command. Output fields are listed in the approximate order in which they appear.

**Table 29: clear services nat mappings app Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of a services interface.
<b>Service set</b>	Name of the service set from which flows are being cleared.
<b>Mappings removed</b>	Number of mappings removed.
<b>Flows removed</b>	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

<b>Syntax</b>	clear services nat mappings eim <b4address <i>b4address/prefix</i> > <subscriber <i>subscriber-ipv4-address</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 14.1.
<b>Description</b>	Clear endpoint independent (EIM) and port control protocol (PCP) mappings .
<b>Options</b>	<p><b>none</b>—Clear all EIM and PCP mappings.</p> <p><b>b4address <i>b4address/prefix</i></b>—(Optional) Clear EIM and PCP mappings for a particular subscriber <i>b4address/prefix</i></p> <p><b>internal-host <i>ipv4address/prefix</i></b>—(Optional) Clear EIM and PCP mappings matching the specified <i>b4address</i> and <i>internal-host</i>..</p> <p><b>port <i>port</i></b>—(Optional) Clear EIM and PCP mappings matching the specified <i>b4address</i>, <i>internal host</i>, and <i>port</i>.</p> <p><b>service-set <i>service-set</i></b>—(Optional) Clear EIM and PCP mappings for the specified <i>service set</i>.</p> <p><b>subscriber <i>subscriber-ipv4-address/prefix</i></b>—(Optional) Clear EIM and PCP mappings for a particular subscriber <i>ipv4-address/prefix</i></p> <ul style="list-style-type: none"> <li><b>port <i>port</i></b>—(Optional) Clear EIM and PCP mappings matching the specified <i>ipv4-address/prefix</i> and <i>port</i>.</li> <li><b>service-set <i>service-set</i></b>—(Optional) Clear EIM and PCP mappings for the specified <i>service set</i>.</li> </ul>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">show services nat mappings on page 1171</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear services nat mappings eim on page 992</a>
<b>Output Fields</b>	Table 30 on page 991 lists the output fields for the <b>clear services nat mappings eim</b> command. Output fields are listed in the approximate order in which they appear.

Table 30: 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 30: 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 pcp

<b>Syntax</b>	clear services nat mappings pcp <b4address <i>b4address/prefix</i> > <subscriber <i>subscriber-ipv4-address</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 14.1.
<b>Description</b>	Clear NAT mappings for Port Control Protocol (PCP).
<b>Options</b>	<p><b>none</b>—Clear all NAT PCP mappings.</p> <p><b>b4address <i>b4address/prefix</i></b>—(Optional) Clear NAT PCP mappings for a particular subscriber <i>b4address/prefix</i></p> <p><b>port <i>port</i></b>—(Optional) Clear NAT PCP mappings matching the specified <i>b4address</i> internal host, and port.</p> <p><b>service-set <i>service-set</i></b>—(Optional) Clear NAT PCP mappings for the specified service set.</p> <p><b>subscriber <i>ipv4-address/prefix</i></b>—(Optional) Clear NAT PCP mappings for a particular subscriber <i>ipv4-address/prefix</i></p> <p><b>port <i>port</i></b>—(Optional) Clear NAT PCP mappings matching the specified <i>ipv4-address/prefix</i>, and port.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show services nat mappings on page 1171</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear services nat mappings pcp on page 994</a>
<b>Output Fields</b>	Table 31 on page 993 lists the output fields for the <b>clear services nat mappings pcp</b> command. Output fields are listed in the approximate order in which they appear.

**Table 31: clear services nat mappings pcp 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 pcp

```
user@host> clear services nat mappings pcp
Interface  Service set  Mappings removed  Flows removed
sp-2/0/0   ss1          0                 0
```

## clear security pki ca-certificate

---

<b>Syntax</b>	clear security pki ca-certificate (all   ca-profile <i>ca-profile-name</i> )
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	Delete certificate authority (CA) digital certificates from the router.
<b>Options</b>	<p><b>all</b>—Delete all CA digital certificates from the router.</p> <p><b>ca-profile <i>ca-profile-name</i></b>—Delete the specified CA profile.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">request security pki ca-certificate enroll on page 1016</a></li> <li>• <a href="#">request security pki ca-certificate load on page 1017</a></li> <li>• <a href="#">show security pki ca-certificate on page 1082</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear security pki ca-certificate all on page 995</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	clear security pki certificate-request (all   certificate-id <i>certificate-id-name</i> )
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	Delete manually generated local digital certificate requests from the router.
<b>Options</b>	<p><b>all</b>—Delete all local digital certificate requests from the router.</p> <p><b>certificate-id <i>certificate-id-name</i></b>—Delete the specified local digital certificate and corresponding public/private key pair.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show security pki certificate-request on page 1086</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear security pki certificate-request all on page 996</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	clear security pki crt (all   ca-profile <i>ca-profile-name</i> )
<b>Release Information</b>	Command introduced in Junos 8.1
<b>Description</b>	Delete certificate revocation lists (CRLs) from the router.
<b>Options</b>	<b>all</b> —Delete all CRLs from the router.  <b>ca-profile <i>ca-profile-name</i></b> —Delete CRLs associated with the specified CA profile.
<b>Required Privilege Level</b>	clear
<b>List of Sample Output</b>	<a href="#">clear security pki crt ca-profile all on page 997</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	clear security pki key-pair (all   certificate-id <i>certificate-id-name</i> )
<b>Release Information</b>	Command introduced in Junos OS Release 8.5.
<b>Description</b>	Clear public key infrastructure (PKI) key pair information for local digital certificates from the router.
<b>Options</b>	<p><b>all</b>—Delete all local digital certificates, certificate requests, and the corresponding public and private key pairs from the router.</p> <p><b>certificate-id <i>certificate-id-name</i></b>—Delete the specified local digital certificate and corresponding public/private key pair.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">request security pki local-certificate enroll on page 1023</a></li><li>• <a href="#">show security pki local-certificate on page 1090</a></li></ul>
<b>Output Fields</b>	This command produces no output.

## Sample Output

```
user@host> clear security pki key pair
```



## clear security pki local-certificate

---

<b>Syntax</b>	clear security pki local-certificate <all   certificate-id <i>certificate-id-name</i>   system-generated>
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	Delete local digital certificates, certificate requests, and the corresponding public/private key pairs from the router.
<b>Options</b>	<p><b>all</b>—(Optional) Delete all local digital certificates, certificate requests, and the corresponding public and private key pairs from the router.</p> <p><b>certificate-id <i>certificate-id-name</i></b>—(Optional) Delete the specified local digital certificate and corresponding public and private key pair.</p> <p><b>system-generated</b>—(Optional) Auto-generated self-signed certificate.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">request security pki local-certificate enroll on page 1023</a></li> <li>• <a href="#">show security pki local-certificate on page 1090</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear security pki local-certificate all on page 999</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	clear services service-sets statistics integrity-drops <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 13.3
<b>Description</b>	Clear integrity-drops statistics for one adaptive services interface, for all adaptive services interfaces, or for one service-set.
<b>Options</b>	<p><b>none</b>—Clear integrity-drops statistics for all configured adaptive service interfaces/ service-set.</p> <p><b>Service-set <i>service-set-name</i></b> —(Optional) Clear integrity-drops statistics for the specified service-set</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear integrity-drops statistics for the specified adaptive services interface.</p>
<b>Required Privilege Level</b>	network
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show services service-sets statistics packet-drops on page 1192</a></li><li>• </li></ul>

## clear services service-sets statistics packet-drops

<b>Syntax</b>	clear services service-sets statistics packet-drops <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 7.4.
<b>Description</b>	Clear dropped-packet statistics for one adaptive services interface or for all adaptive services interfaces.
<b>Options</b>	<p><b>none</b>—Clear dropped-packet statistics for all configured adaptive services interfaces.</p> <p><b>interface <i>interface-name</i></b>—(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>.</p>
<b>Required Privilege Level</b>	network
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show services service-sets statistics packet-drops on page 1192</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear services service-sets statistics packet-drops on page 1001</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	<code>clear services service-sets statistics syslog</code> <code>&lt;service-set <i>service-set-name</i>&gt;</code> <code>&lt;interface <i>interface-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1.
<b>Description</b>	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.
<b>Options</b>	<b>none</b> —Clear system log for all configured services interfaces and their service sets.  <b>interface <i>interface-name</i></b> —(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> .  <b>service-set <i>service-set-name</i></b> —(Optional) Clear system log statistics for the specified services interface.
<b>Required Privilege Level</b>	network
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show services service-sets statistics syslog on page 1194</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear services service-sets statistics syslog on page 1002</a>
<b>Output Fields</b>	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 sessions

<b>Syntax</b>	<pre>clear services sessions &lt;application-protocol <i>protocol</i>&gt; &lt;destination-port <i>destination-port</i>&gt; &lt;destination-prefix <i>destination-prefix</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;ip-action&gt; &lt;protocol <i>protocol</i>&gt; &lt;service-set <i>service-set</i>&gt; &lt;source-port <i>source-port</i>&gt; &lt;source-prefix <i>source-prefix</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 13.1.</p> <p>Command introduced in Junos OS Service Control Gateway Release 14.1X55.</p>
<b>Description</b>	<p>Clear services sessions currently active on the embedded Junos PIC/MIC. When you enter this command, the sessions are marked for deletion and are cleared thereafter. The time that is taken to clear the currently active sessions varies, depending on the scaled nature of the environment.</p>
<b>Options</b>	<p><b>none</b>—Clear all sessions.</p> <p><b>application-protocol</b>—(Optional) Clear sessions for one of the following application protocols:</p> <ul style="list-style-type: none"> <li>• <b>bootp</b>—Bootstrap protocol</li> <li>• <b>dce-rpc</b>—Distributed Computing Environment-Remote Procedure Call protocols</li> <li>• <b>dce-rpc-portmap</b>—Distributed Computing Environment-Remote Procedure Call protocols portmap service</li> <li>• <b>dns</b>—Domain Name System protocol</li> <li>• <b>exec</b>—Exec</li> <li>• <b>ftp</b>—File Transfer Protocol</li> <li>• <b>h323</b>—H.323 standards</li> <li>• <b>icmp</b>—Internet Control Message Protocol</li> <li>• <b>iiop</b>—Internet Inter-ORB Protocol</li> <li>• <b>ip</b>—IP</li> <li>• <b>login</b>—Login</li> <li>• <b>netbios</b>—NetBIOS</li> <li>• <b>netshow</b>—NetShow</li> <li>• <b>pptp</b>—Point-to-Point Tunneling Protocol</li> <li>• <b>realaudio</b>—RealAudio</li> <li>• <b>rpc</b>—Remote Procedure Call protocol</li> </ul>

- **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

**destination-port** *destination-port*—(Optional) Clear sessions for a particular destination port. The range of values is from 0 to 65535.

**destination-prefix** *destination-prefix*—(Optional) Clear sessions for a particular destination prefix.

**interface** *interface-name*—(Optional) Clear sessions for a particular interface. On M Series and T Series routers, the *interface-name* can be **ms-fpc/pic/port** or **rspnumber**.

**ip-action**—(Optional) 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 policy-name rulebase-ips rule rule-name then}** hierarchy level.

**protocol**—(Optional) Clear sessions 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
- **icmp6**—Internet Control Message Protocol version 6
- **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 Transmission Protocol

- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

**service-set** *service-set*—(Optional) Clear sessions for a particular service set.

**source-port** *source-port*—(Optional) Clear sessions for a particular source port. The range of values is from 0 through 65535.

**source-prefix** *source-prefix*—(Optional) Clear sessions for a particular source prefix.

**Required Privilege Level**

clear

**Related Documentation**

- [show services sessions on page 1202](#)

**List of Sample Output** [clear services sessions on page 1005](#)

**Output Fields** [Table 32 on page 1005](#) lists the output fields for the **clear services sessions** command. Output fields are listed in the approximate order in which they appear.

**Table 32: clear services sessions Output Fields**

Field Name	Field Description
Interface	Name of an interface.
Service set	Name of the service set from which sessions are being cleared.
Sessions marked for deletion	Number of sessions that are marked for deletion and are subsequently cleared.

## Sample Output

### clear services sessions

```

user@hosts>clear services sessions
Interface  Service set      Sessions marked
for deletion
ms-0/0/0   sset              10

```

## 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 1228](#)

**List of Sample Output** [clear services stateful-firewall flows on page 1007](#)

**Output Fields** [Table 33 on page 1007](#) 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 33: clear services stateful-firewall flows Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of an adaptive services interface.
<b>Service set</b>	Name of the service set from which flows are being cleared.
<b>Conv removed</b>	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 1234](#)

**List of Sample Output** [clear services stateful-firewall sip-call on page 1010](#)

**Output Fields** [Table 34 on page 1010](#) 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 34: 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 1239](#)

List of Sample Output

[clear services stateful-firewall sip-register on page 1013](#)

**Output Fields** Table 35 on page 1013 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 35: 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><b>none</b>—Clear stateful firewall statistics for all interfaces and all service sets.</p> <p><b>interface <i>interface-name</i></b>—(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><b>service-set <i>service-set</i></b>—(Optional) Clear stateful firewall statistics for the specified service set.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">show services stateful-firewall statistics on page 1243</a></li></ul>
List of Sample Output	<a href="#">clear services stateful-firewall statistics on page 1014</a>
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)

<b>Syntax</b>	request interface (revert   switchover) ( <i>rspnumber</i>   <i>rlsnumber</i> )
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Support for <b>rlsq</b> interfaces added in Junos OS Release 7.6.
<b>Description</b>	(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>  <b>NOTE:</b> All <b>rlsq</b> switchover or revert operations are allowed from the <b>rlsnumber</b> level only and not for individual channelized interfaces (<b>rlsnumber:unit</b>). </div>	
<p>On an aggregated Ethernet interface with link protection enabled, use the <b>request interface (revert   switchover)</b> (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>	
<b>Options</b>	<p><b>(revert   switchover)</b>—The <b>revert</b> keyword restores active processing to the primary adaptive services (<b>sp</b>) or link services IQ (<b>lsq</b>) interface. The <b>switchover</b> keyword transfers active processing to the secondary (backup) interface.</p> <p><b>rspnumber</b>—Redundant adaptive services interface name.</p> <p><b>rlsnumber</b>—Redundant link services IQ interface name.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">request interface revert on page 1015</a> <a href="#">request interface switchover on page 1015</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	request security pki ca-certificate enroll ca-profile <i>ca-profile-name</i>
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	Request a digital certificate from a certificate authority (CA) online by using the Simple Certificate Enrollment Protocol (SCEP).
<b>Options</b>	<b>ca-profile <i>ca-profile-name</i></b> —CA profile name.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">clear security pki ca-certificate on page 995</a></li><li>• <a href="#">show security pki ca-certificate on page 1082</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request security pki ca-certificate enroll on page 1016</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	<code>request security pki ca-certificate load ca-profile <i>ca-profile-name</i> filename <i>path/filename</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	Manually load a certificate authority (CA) digital certificate from a specified location.
<b>Options</b>	<p><b>ca-profile <i>ca-profile-name</i></b>—Load the specified CA profile.</p> <p><b>filename <i>path/filename</i></b>—Directory location and filename of the CA digital certificate.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear security pki ca-certificate on page 995</a></li> <li>• <a href="#">show security pki ca-certificate on page 1082</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request security pki ca-certificate load on page 1017</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	<code>request security pki ca-certificate verify ca-profile <i>ca-profile-name</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 8.5.
<b>Description</b>	Verify the digital certificate installed for the specified certificate authority (CA).
<b>Options</b>	<b>ca-profile <i>ca-profile-name</i></b> —Name of the local digital certificate identifier.
<b>Required Privilege Level</b>	maintenance
<b>Output Fields</b>	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

---

<b>Syntax</b>	<code>request security pki crt load ca-profile <i>ca-profile-name</i> filename <i>path/filename</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 8.1.
<b>Description</b>	Manually install a certificate revocation list (CRL) on the router from a specified location.
<b>Options</b>	<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.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">request security pki crt load on page 1019</a>
<b>Output Fields</b>	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

---

<b>Syntax</b>	<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> &lt;email <i>email-address</i>&gt; &lt;filename (<i>path</i>   <i>terminal</i>)&gt; &lt;ip-address <i>ip-address</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	Manually generate a local digital certificate request in the Public-Key Cryptography Standards #10 (PKCS-10) format.
<b>Options</b>	<p><b>certificate-id <i>certificate-id-name</i></b>—Name of the local digital certificate and the public/private key pair.</p> <p><b>domain-name <i>domain-name</i></b>—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><b>subject <i>subject-distinguished-name</i></b>—Distinguished name format that contains the common name, department, company name, state, and country:</p> <ul style="list-style-type: none"><li>• <b>CN</b>—Common name</li><li>• <b>OU</b>—Organizational unit name</li><li>• <b>O</b>—Organization name</li><li>• <b>ST</b>—State</li><li>• <b>C</b>—Country</li></ul> <p><b>email <i>email-address</i></b>—(Optional) E-mail address of the certificate holder.</p> <p><b>filename (<i>path</i>   <i>terminal</i>)</b>—(Optional) Location where the local digital certificate request should be placed or the login terminal.</p> <p><b>ip-address <i>ip-address</i></b>—(Optional) IP address of the router.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">clear security pki certificate-request on page 996</a></li><li>• <a href="#">show security pki certificate-request on page 1086</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request security pki generate-certificate-request on page 1021</a>
<b>Output Fields</b>	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.example.net filename entrust-req2 subject cn=router2.example.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

---

<b>Syntax</b>	<code>request security pki generate-key-pair certificate-id <i>certificate-id-name</i> &lt;size (512   1024   2048)&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	Generate a Public Key Infrastructure (PKI) public and private key pair for a local digital certificate.
<b>Options</b>	<b>certificate-id</b> <i>certificate-id-name</i> —Name of the local digital certificate and the public/private key pair.  <b>size</b> —(Optional) Key pair size. The key pair size can be <b>512</b> , <b>1024</b> , or <b>2048</b> bits.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">request security pki generate-key-pair on page 1022</a>
<b>Output Fields</b>	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

<b>Syntax</b>	request security pki local-certificate enroll <i>ca-profile ca-profile-name</i> certificate-id <i>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> >
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	Request that a certificate authority (CA) enroll and install a local digital certificate online by using the Simple Certificate Enrollment Protocol (SCEP).
<b>Options</b>	<p><b>ca-profile</b> <i>ca-profile-name</i>—CA profile name.</p> <p><b>certificate-id</b> <i>certificate-id-name</i>—Name of the local digital certificate and the public/private key pair.</p> <p><b>challenge-password</b> <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><b>domain-name</b> <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><b>subject</b> <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"> <li>• <b>CN</b>—Common name</li> <li>• <b>OU</b>—Organizational unit name</li> <li>• <b>O</b>—Organization name</li> <li>• <b>ST</b>—State</li> <li>• <b>C</b>—Country</li> </ul> <p><b>email</b> <i>email-address</i>—(Optional) E-mail address of the certificate holder.</p> <p><b>ip-address</b> <i>ip-address</i>—(Optional) IP address of the router.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show security pki local-certificate on page 1090</a></li> </ul>
<b>Output Fields</b>	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.example.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

<b>Syntax</b>	<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>
<b>Release Information</b>	Command introduced in Junos OS Release 9.1.
<b>Description</b>	Manually generate a self-signed certificate for the given distinguished name.
<b>Options</b>	<p><b>certificate-id <i>certificate-id-name</i></b>—Name of the local digital certificate and the public/private key pair.</p> <p><b>domain-name <i>domain-name</i></b>—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><b>email <i>email-address</i></b>—E-mail address of the certificate holder.</p> <p><b>ip-address <i>ip-address</i></b>—IP address of the router.</p> <p><b>subject <i>subject-distinguished-name</i></b>—Distinguished name format that contains the common name, department, company name, state, and country:</p> <ul style="list-style-type: none"> <li>• <b>CN</b>—Common name</li> <li>• <b>OU</b>—Organizational unit name</li> <li>• <b>O</b>—Organization name</li> <li>• <b>ST</b>—State</li> <li>• <b>C</b>—Country</li> </ul>
<b>Required Privilege Level</b>	<p><b>maintenance</b></p> <p><b>security</b></p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show security pki local-certificate (View)</a></li> </ul>
<b>Output Fields</b>	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 example.net email user1@example.net
Self-signed certificate generated and loaded successfully
```

## request security pki local-certificate load

---

<b>Syntax</b>	<code>request security pki local-certificate load certificate-id <i>certificate-id-name</i> filename <i>path</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	Manually load a local digital certificate from a specified location.
<b>Options</b>	<p><b>certificate-id <i>certificate-id-name</i></b>—Name of the public/private key pair mapped to the local digital certificate.</p> <p><b>filename <i>path/filename</i></b>—Directory location and filename of the local digital certificate provided by the CA.</p>
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">request security pki local-certificate load on page 1026</a>
<b>Output Fields</b>	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

<b>Syntax</b>	<code>request security pki local-certificate verify certificate-id <i>certificate-id-name</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 8.5.
<b>Description</b>	Verify the validity of the local digital certificate identifier.
<b>Options</b>	<code>certificate-id <i>certificate-id-name</i></code> —Display the specified certificate identifier name.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show security pki local-certificate on page 1090</a></li> </ul>
<b>Output Fields</b>	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

---

<b>Syntax</b>	<code>request services ipsec-vpn ipsec switch tunnel local-gateway <i>address</i> remote-gateway <i>address</i></code> <code>&lt;routing-instance <i>instance-name</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. <b>routing-instance</b> option added in Release 8.1.
<b>Description</b>	(Adaptive services interface only) Manually switch between primary and backup IP Security (IPsec) tunnels.
<b>Options</b>	<b>local-gateway <i>address</i></b> —Gateway address of the local system.  <b>remote-gateway <i>address</i></b> —Gateway address of the remote system.  <b>routing-instance <i>instance-name</i></b> —(Optional) VRF instance associated with local gateway address.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show services ipsec-vpn ipsec security-associations on page 1124</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request services ipsec-vpn ipsec switch tunnel on page 1028</a>
<b>Output Fields</b>	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)

<b>Syntax</b>	<pre>show interfaces <i>interface-type</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display status information about the specified adaptive services interface.
<b>Options</b>	<p><b><i>interface-type</i></b>—On M Series and T Series routers, the interface type is <b>sp-<i>fpc/pic/port</i></b>.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show interfaces (Adaptive Services) on page 1034</a></p> <p><a href="#">show interfaces brief (Adaptive Services) on page 1034</a></p> <p><a href="#">show interfaces detail (Adaptive Services) on page 1034</a></p> <p><a href="#">show interfaces extensive (Adaptive Services) on page 1035</a></p>
<b>Output Fields</b>	Table 36 on page 1029 lists the output fields for the <b>show interfaces</b> (adaptive services and redundant adaptive services) command. Output fields are listed in the approximate order in which they appear.

**Table 36: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields**

Field Name	Field Description	Level of Output
Physical Interface		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Physical interface's index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>

**Table 36: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields**  
(continued)

Field Name	Field Description	Level of Output
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Type</b>	Encapsulation being used on the interface.	All levels
<b>Link-level type</b>	Encapsulation being used on the physical interface.	All levels
<b>MTU</b>	MTU size on the physical interface.	All levels
<b>Clocking</b>	Reference clock source: can be <b>Internal</b> or <b>External</b> .	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Device flags</b>	Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Interface flags</b>	Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Link type</b>	Physical interface link type: <b>Full-Duplex</b> or <b>Half-Duplex</b> .	<b>detail extensive none</b>
<b>Link flags</b>	Information about the link. Possible values are described in the "Link Flags" section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Physical info</b>	Information about the physical interface.	<b>detail extensive</b>
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>
<b>Current address</b>	Configured MAC address.	<b>detail extensive none</b>
<b>Hardware address</b>	MAC address of the hardware.	<b>detail extensive none</b>
<b>Alternate link address</b>	Backup address of the link.	<b>detail extensive none</b>
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .	<b>detail extensive none</b>
<b>Input Rate</b>	Input rate in bits per second (bps) and packets per second (pps).	None specified
<b>Output Rate</b>	Output rate in bps and pps.	None specified
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>



Table 36: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Input errors</b>	<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"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Frames received smaller than the runt threshold.</li> <li>• <b>Giants</b>—Frames received larger than the giant threshold.</li> <li>• <b>Policed discards</b>—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.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Output errors</b>	<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"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. 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.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>MTU errors</b>—Number of packets larger than the MTU threshold.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Logical interface index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP interface index number.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

**Table 36: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (continued)**

Field Name	Field Description	Level of Output
<b>Flags</b>	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Encapsulation</b>	Encapsulation on the logical interface.	All levels
<b>Input packets</b>	Number of packets received on the logical interface.	None specified
<b>Output packets</b>	Number of packets transmitted on the logical interface.	None specified
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the logical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Local statistics</b>	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.	<b>detail extensive</b>
<b>Transit statistics</b>	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.	<b>detail extensive</b>
<b><i>protocol-family</i></b>	Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.	<b>brief</b>
<b>Protocol</b>	Protocol family configured on the logical interface, such as <b>iso</b> , <b>inet6</b> , <b>mpls</b> .	<b>detail extensive none</b>
<b>MTU</b>	MTU size on the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route table</b>	Routing table in which the logical interface address is located. For example, <b>0</b> refers to the routing table <b>inet.0</b> .	<b>detail extensive</b>
<b>Flags</b>	Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Addresses, Flags</b>	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>

**Table 36: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields**  
(continued)

Field Name	Field Description	Level of Output
<b>Broadcast</b>	Broadcast address.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

## 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)

<b>Syntax</b>	<pre>show interfaces lsq-fpc/pic/port &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;l2-statistics&gt; &lt;media&gt; &lt;snmp-index snmp-index&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p><b>l2-statistics</b> option introduced with Junos OS Release 12.1.</p>
<b>Description</b>	(M Series, MX Series, and T Series routers only) Display status information about the specified link services intelligent queuing (IQ) interface.
<b>Options</b>	<p><b>lsq-fpc/pic/port</b>—Display standard status information about the specified link services IQ interface.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>l2-statistics</b>—(Optional) Display Layer 2 queue statistics for Multilink Point-to-Point Protocol (MLPPP), FRF.15, and FRF.16 bundles.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index snmp-index</b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Additional Information</b>	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.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Link and Multilink Services Overview</a></li> <li>• <a href="#">Multilink Interfaces on Channelized MICs Overview</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show interfaces extensive (MLPPP on Link Services IQ) on page 1052</a></p> <p><a href="#">show interfaces extensive (Multiclass MLPPP on Link Services IQ) on page 1053</a></p> <p><a href="#">show interfaces extensive (MLPPP on Link Services IQ Bundle) on page 1055</a></p> <p><a href="#">show interfaces extensive (MFR on Link Services IQ Bundle) on page 1056</a></p> <p><a href="#">show interfaces extensive (Multiclass MLPPP on Link Services IQ) on page 1058</a></p>

**Output Fields** Table 37 on page 1038 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 37: show interfaces (Link Services IQ) Output Fields**

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Physical interface index number, which reflects its initialization sequence.	<b>detail extensive</b> none
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive</b> none
<b>Generation</b>	Unique number for use by Juniper Networks technical support.	<b>detail extensive</b>
<b>Link-level type</b>	Encapsulation being used on the physical interface: <b>Multilink-Frame-Relay-UNI-NNI Multilink-Frame-Relay-UNI-NNI</b> (default), <b>LinkService</b> , <b>Frame-relay</b> , <b>Frame-relay-ccc</b> , or <b>Frame-relay-tcc</b> .	All levels
<b>MTU</b>	Maximum transmission unit size on the physical interface.	All levels
<b>Device flags</b>	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface flags</b>	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .	<b>detail extensive</b> none



Table 37: show interfaces (Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Multilink Frame Relay UNI NNI bundle options</b>	<p>(Multilink Frame Relay UNI NNI only) Configured information about Multilink Frame Relay bundle options.</p> <ul style="list-style-type: none"> <li>• <b>Device type</b>—DCE (data communication equipment) or DTE (data terminal equipment).</li> <li>• <b>MRRU</b>—Configured size of the maximum received reconstructed unit (MRRU): 1500 to 4500 bytes. The default is 1524 bytes.</li> <li>• <b>Bandwidth</b>—Speed at which the interface is running.</li> <li>• <b>Fragmentation threshold</b>—Configured fragmentation threshold: 128 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation.</li> <li>• <b>Red differential delay limit</b>—Red differential delay limit among bundle links has been reached, indicating an action will occur.</li> <li>• <b>Yellow differential delay limit</b>—Yellow differential delay among bundle links has been reached, indicating a warning will occur.</li> <li>• <b>Red differential delay action</b>—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>.</li> <li>• <b>Link layer overhead</b>—Percentage of bundle bandwidth to be set aside for link layer overhead.</li> <li>• <b>Reassembly drop timer</b>—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.</li> <li>• <b>Links needed to sustain bundle</b>—Minimum number of links to sustain the bundle: 1 through 8.</li> <li>• <b>LIP Hello timer</b>—Link Interleaving Protocol hello timer: 1 through 180 seconds. <ul style="list-style-type: none"> <li>• <b>Acknowledgement timer</b>—Maximum period to wait for an add link acknowledgement, hello acknowledgement, or remove link acknowledgement: 1 through 10 seconds.</li> <li>• <b>Acknowledgement retries</b>—Number of retransmission attempts to be made for consecutive hello or remove link messages after the expiration of the acknowledgement timer: 1 through 5.</li> </ul> </li> </ul>	<b>detail extensive none</b>

Table 37: 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"> <li>• <b>Bundle class</b>—Bundle class ID.</li> <li>• <b>LMI type</b>—Multilink Frame Relay UNI NNI LMI type: <b>ANSI, Q.933 ANNEX A</b>, or <b>Consortium</b>. <ul style="list-style-type: none"> <li>• <b>T391 LIV polling timer</b>—Multilink Frame Relay UNI NNI Full status polling counter: 1 through 255, with a default value of 6.</li> <li>• <b>T392 polling verification timer</b>—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.</li> <li>• <b>N391 full status polling count</b>—Multilink Frame Relay UNI NNI Full status polling counter: 1 through 255.</li> <li>• <b>N392 error threshold</b>—Multilink Frame Relay UNI NNI LMI error threshold: 1 through 10.</li> </ul> </li> <li>• <b>N393 monitored event count</b>—Multilink Frame Relay UNI NNI LMI monitored event count: 1 through 10, with a default value of 4.</li> <li>• <b>Consortium LMI Settings</b> <ul style="list-style-type: none"> <li>• <b>n391dte</b>—DTE full status polling interval in seconds: 1 through 255.</li> <li>• <b>n392dce</b>—DCE error threshold: 1 through 10.</li> <li>• <b>n392dte</b>—DTE error threshold: 1 through 10.</li> <li>• <b>n393dce</b>—DCE monitored event count: 1 through 10.</li> <li>• <b>n393dte</b>—DTE monitored event count: 1 through 10.</li> <li>• <b>t391dte</b>—DTE polling verification timer (in seconds): 5 through 30.</li> <li>• <b>t392dce</b>—DCE polling verification timer (in seconds): 5 through 30.</li> </ul> </li> </ul>	detail extensive none
LMI	<p>Local Management Interface packet statistics:</p> <ul style="list-style-type: none"> <li>• <b>Input</b>—Number of packets arriving on the interface (nn) and timestamp of the most recent packet arrival, in the format: <b>Input: nn (last seen hh:mm:ss ago)</b></li> <li>• <b>Output</b>—Number of packets sent out on the interface (nn) and how much time has passed since the last packet was sent, in the format: <b>Output: nn (last seen hh:mm:ss ago)</b></li> </ul>	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"> <li>• <b>Enquiries sent</b>—Number of link status enquiries sent from the DTE to the DCE.</li> <li>• <b>Full enquiries sent</b>—Number of full enquiries sent from the DTE to the DCE.</li> <li>• <b>Enquiry responses received</b>—Number of enquiry responses received by the DCE from the DTE.</li> <li>• <b>Full enquiry responses received</b>—Number of full enquiry responses received by DCE from the DTE.</li> </ul>	detail extensive none

Table 37: show interfaces (Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>DCE Statistics</b>	<p>Statistics about information transferred from the DCE to the DTE.</p> <ul style="list-style-type: none"> <li>• <b>Enquiries received</b>—Number of enquiries received by the DCE from the DTE.</li> <li>• <b>Full enquiries received</b>—Number of full enquiries received by the DCE from the DTE.</li> <li>• <b>Enquiry responses sent</b>—Number of enquiry responses sent from the DCE to the DTE.</li> <li>• <b>Full enquiry responses sent</b>—Number of full enquiry responses sent from the DCE to the DTE.</li> </ul>	<b>detail extensive none</b>
<b>Common Statistics</b>	<p>Statistics about messages snet between the DTE and the DCE.</p> <ul style="list-style-type: none"> <li>• <b>Unknown messages received</b>—Number of received packets that do not fall into any other category.</li> <li>• <b>Asynchronouts updates received</b>—Number of link status peer changes received.</li> <li>• <b>Out-of-sequence packets received</b>—Number of packets for which the sequence of the packets received is different from the expected sequence.</li> <li>• <b>Keepalive responses timed out</b>—Number of keepalive responses that time out when no Local Management Interface (LMI) packet was reported for <b>n392dte</b> or <b>n393dce</b> intervals. (See <i>LMI settings</i>.)</li> </ul>	
<b>Traffic statistics</b>	<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>	<b>detail extensive</b>
<b>DLCInn</b>	<p>Data-link connection identifier (DLCI) number of the logical interface. The following information is displayed.</p> <ul style="list-style-type: none"> <li>• <b>Flags</b>—Values are: <ul style="list-style-type: none"> <li>• <b>Active</b>—Set when the link is active and the DTE and DCE are exchanging information.</li> <li>• <b>Down</b>—Set when the link is active, but no information is received from the DTE.</li> <li>• <b>DCE unconfigured</b>—Set when the corresponding DLCI in the DCE is not configured.</li> <li>• <b>Configured</b>—Set when the corresponding DLCCI is configured.</li> <li>• <b>DCE-Configured</b>—Displayed when the command is issued from the DTE.</li> </ul> </li> </ul>	
<b>DLCI Statistics</b>	<p>(Frame Relay) Data-link connection identifier (DLCI) statistics.</p> <ul style="list-style-type: none"> <li>• <b>Active DLCI</b>—Number of active DLCIs.</li> <li>• <b>Inactive DLCI</b>—Number of inactive DLCIs.</li> </ul>	
<b>Input rate</b>	(Redundant LSQ) Rate of bits and packets received on the interface.	None specified
<b>Output rate</b>	(Redundant LSQ) Rate of bits and packets transmitted on the interface.	None specified

Table 37: show interfaces (Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>
<b>Traffic statistics</b>	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.	<b>detail extensive</b>
<b>Frame exceptions</b>	<p>Information about framing exceptions. Includes events recorded under <b>Exception Events</b> for each logical interface.</p> <ul style="list-style-type: none"> <li>• <b>Oversized frames</b>—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb (kilobits).</li> <li>• <b>Errored input frames</b>—Number of input frame errors.</li> <li>• <b>Input on disabled link/bundle</b>—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.</li> <li>• <b>Output for disabled link/bundle</b>—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.</li> <li>• <b>Queuing drops</b>—Total number of packets dropped before traffic enters the link services IQ interface. Indicates that the interface is becoming oversubscribed.</li> </ul>	<b>extensive</b>
<b>Buffering exceptions</b>	<p>Information about buffering exceptions. Includes events recorded under <b>Exception Events</b> for each logical interface:</p> <ul style="list-style-type: none"> <li>• <b>Packet data buffer overflow</b>—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity.</li> <li>• <b>Fragment data buffer overflow</b>—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.</li> </ul>	<b>extensive</b>

Table 37: show interfaces (Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Assembly exceptions</b>	<p>(Multilink Frame Relay end-to-end only) Information about assembly exceptions. Includes events recorded under <b>Exception Events</b> 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"> <li>• <b>Fragment timeout</b>—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.</li> <li>• <b>Missing sequence number</b>—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.</li> <li>• <b>Out-of-order sequence number</b>—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.</li> <li>• <b>Out-of-range sequence number</b>—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.</li> </ul>	<b>extensive</b>
<b>Hardware errors (sticky)</b>	<p>(Multilink Frame Relay end-to-end only) Information about hardware errors:</p> <ul style="list-style-type: none"> <li>• <b>Data memory error</b>—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support.</li> <li>• <b>Control memory error</b>—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support.</li> </ul>	<b>extensive</b>
<b>Egress queues</b>	Total number of egress queues supported on the specified interface.	<b>detail extensive none</b>

Table 37: show interfaces (Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Queue counters</b>	Queue number and its associated user-configured forwarding class name. <ul style="list-style-type: none"> <li><b>Queued packets</b>—Number of queued packets.</li> <li><b>Transmitted packets</b>—Number of transmitted packets.</li> <li><b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>	<b>detail extensive none</b>
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Logical interface index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	Logical interface SNMP interface index number.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support.	<b>detail extensive</b>
<b>Flags</b>	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Encapsulation</b>	Encapsulation being used: PPP or Multilink PPP.	All levels
<b>Bandwidth</b>	Speed at which the interface is running.	All levels
<b>Bundle options</b>	(Multilink Frame Relay end-to-end interfaces only) <ul style="list-style-type: none"> <li><b>MRRU</b>—Configured size of the maximum received reconstructed unit (MRRU): 1500 through 4500 bytes. The default is 1504 bytes.</li> <li><b>Drop timer period</b>—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.</li> <li><b>Sequence number format</b>—Short sequence number header format (MLPPP only).</li> <li><b>Fragmentation threshold</b>—Configured fragmentation threshold: 64 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation.</li> <li><b>Links needed to sustain bundle</b>—Minimum number of links to sustain the bundle: 1 through 8.</li> <li><b>Multilink classes</b>—Number of multilink classes negotiated.</li> <li><b>Link layer overhead</b>—Percentage of bundle bandwidth to be set aside for link-layer overhead.</li> </ul>	<b>detail extensive none</b>

Table 37: show interfaces (Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Bundle status</b> (MLPPP) or <b>Multilink class status</b> (Multiclass MLPPP)	Information about bundle status: <ul style="list-style-type: none"> <li>• <b>Remote MRRU</b>—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed.</li> <li>• <b>Received sequence number</b>—Sequence number for received packets.</li> <li>• <b>Transmitted sequence number</b>—Sequence number for transmitted packets.</li> <li>• <b>Packet drops</b>—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.</li> <li>• <b>Fragment drops</b>—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.</li> <li>• <b>MRRU exceeded</b>—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release.</li> <li>• <b>Fragment timeout</b>—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.</li> <li>• <b>Missing sequence number</b>—A gap was detected in the sequence numbers of fragments on a bundle.</li> <li>• <b>Out-of-order sequence number</b>—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.</li> <li>• <b>Out-of-range sequence number</b>—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.</li> <li>• <b>Packet data buffer overflow</b>—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity.</li> <li>• <b>Fragment data buffer overflow</b>—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.</li> </ul>	<b>detail extensive none</b>

Table 37: show interfaces (Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Statistics</b>	<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"> <li>• <b>Bundle</b>—Information for each active bundle link. <ul style="list-style-type: none"> <li>• <b>Fragments: Input and Output</b>—Total number and rate of fragments received and transmitted.</li> <li>• <b>Packets: Input and Output</b>—Total number and rate of packets received and transmitted.</li> <li>• <b>Multilink class</b>—(Multiclass MLPPP only) Information about multiclass links used in the multilink operation.</li> </ul> </li> <li>• <b>Link</b>—Information about links used in the multilink operation. <ul style="list-style-type: none"> <li>• <b>Link name</b>—Interface name of the link services IQ channel and state information (physical link <b>up</b> or <b>down</b>).</li> <li>• <b>Input and Output</b>—Total number and rate of fragments and packets received and transmitted.</li> </ul> </li> </ul>	<b>detail extensive</b>
<b>NCP state</b>	<p>(PPP) Network Control Protocol state.</p> <ul style="list-style-type: none"> <li>• <b>Conf-ack-received</b>—Acknowledgement was received.</li> <li>• <b>Conf-ack-sent</b>—Acknowledgement was sent.</li> <li>• <b>Conf-req-sent</b>—Request was sent.</li> <li>• <b>Down</b>—NCP negotiation is incomplete (not yet completed or has failed).</li> <li>• <b>Not-configured</b>—NCP is not configured on the interface.</li> <li>• <b>Opened</b>—NCP negotiation is successful.</li> </ul>	<b>detail extensive none</b>
<b>Protocol</b>	Protocol family configured on the logical interface.	<b>detail extensive none</b>
<b>MTU</b>	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 <b>Adjusted</b> .	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route Table</b>	Routing table in which this address exists. For example, <b>Route table:0</b> refers to inet.0.	<b>detail extensive</b>
<b>Flags</b>	Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Addresses, Flags</b>	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> .	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>



Table 37: show interfaces (Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address on the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support.	<b>detail extensive</b>
<b>MLPPP Bundle Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Logical interface index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	Logical interface SNMP interface index number.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support.	<b>detail extensive</b>
<b>Flags</b>	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>SNMP-Traps</b>	SNMP trap notifications are enabled.	All levels
<b>Encapsulation</b>	Encapsulation being used: PPP, Multilink PPP, or Multilink-FR.	All levels
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: <i>year-month-day hour:minute:second timezone (hour:minute:second ago)</i></b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .	<b>detail extensive none</b>
<b>Bandwidth</b>	Speed at which the interface is running.	All levels
<b>Bundle links information</b>	Information about the bundled links. <ul style="list-style-type: none"> <li>• <b>Active bundle links</b>—Number of active links.</li> <li>• <b>Removed bundle links</b>—Information about links used in the multilink operation.</li> <li>• <b>Disabled bundle links</b>—Number of disabled links.</li> </ul>	<b>detail extensive none</b>

Table 37: show interfaces (Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Bundle options</b>	<p>(Multilink Frame Relay end-to-end interfaces only)</p> <ul style="list-style-type: none"> <li>• <b>MRRU</b>—Configured size of the maximum received reconstructed unit (MRRU): 1500 through 4500 bytes. The default is 1504 bytes.</li> <li>• <b>Drop timer period</b>—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.</li> <li>• <b>Inner PPP Protocol field compression</b>—Inner PPP protocol compression is enabled or disabled.</li> <li>• <b>Sequence number format</b>—Short sequence number header format (MLPPP only).</li> <li>• <b>Fragmentation threshold</b>—Configured fragmentation threshold: 64 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation.</li> <li>• <b>Links needed to sustain bundle</b>—Minimum number of links to sustain the bundle: 1 through 8.</li> <li>• <b>Multilink classes</b>—Number of multilink classes negotiated.</li> <li>• <b>Link layer overhead</b>—Percentage of bundle bandwidth to be set aside for link-layer overhead.</li> </ul>	<b>detail extensive none</b>

Table 37: show interfaces (Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Bundle status</b> (MLPPP)	<p>Information about bundle status:</p> <ul style="list-style-type: none"> <li>• <b>Received sequence number</b>—Sequence number for received packets.</li> <li>• <b>Transmit sequence number</b>—Sequence number for transmitted packets.</li> <li>• <b>Packet drops</b>—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.</li> <li>• <b>Fragment drops</b>—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.</li> <li>• <b>MRRU exceeded</b>—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release.</li> <li>• <b>Fragment timeout</b>—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.</li> <li>• <b>Missing sequence number</b>—A gap was detected in the sequence numbers of fragments on a bundle.</li> <li>• <b>Out-of-order sequence number</b>—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.</li> <li>• <b>Out-of-range sequence number</b>—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.</li> <li>• <b>Packet data buffer overflow</b>—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity.</li> <li>• <b>Fragment data buffer overflow</b>—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.</li> </ul>	<b>detail extensive none</b>

Table 37: show interfaces (Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Statistics</b>	<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"> <li>• Input side - Total overhead = 6 bytes. <ul style="list-style-type: none"> <li>• ML: 4 bytes of ML header = 1 byte of Flag + 3 bytes of long sequence number.</li> <li>• PPP: 2 bytes of protocol field.</li> </ul> </li> <li>• Output side - Total overhead = 11 bytes. <ul style="list-style-type: none"> <li>• ML: 4 bytes of ML Header = 1 byte of Flag + 3 bytes of Long sequence number.</li> <li>• PPP: 5 bytes = 4 bytes of header + 1 byte of Idle flag.</li> <li>• 2 bytes of Service Cookie.</li> </ul> </li> <li>• <b>Bundle</b>—Information for each active bundle link. <ul style="list-style-type: none"> <li>• <b>Multilink: Input and Output</b>—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.</li> <li>• <b>Network: Input and Output</b>—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.</li> </ul> </li> <li>• <b>Link</b>—Information about links used in the multilink operation. <ul style="list-style-type: none"> <li>• <b>Link name</b>—The interface name of the link services IQ channel and state information (physical link <i>up</i> or <i>down</i>) and up time.</li> <li>• <b>Input and Output</b>—Total number and rate of frames, bytes, and bits per second received and transmitted.</li> </ul> </li> </ul>	<b>extensive</b>

Table 37: show interfaces (Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Multilink detail statistics</b>	<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"> <li>• Input side - Total overhead = 6 bytes. <ul style="list-style-type: none"> <li>• ML: 4 bytes of ML header = 1 byte of Flag + 3 bytes of long sequence number.</li> <li>• PPP: 2 bytes of protocol field.</li> </ul> </li> <li>• Output side - Total overhead = 9 bytes. <ul style="list-style-type: none"> <li>• ML: 4 bytes of ML Header = 1 byte of Flag + 3 bytes of Long sequence number.</li> <li>• PPP: 5 bytes = 4 bytes of header + 1 byte of Idle flag.</li> </ul> </li> <li>• <b>Bundle</b>—Information for the bundle link. <ul style="list-style-type: none"> <li>• <b>Fragments: Input and Output</b>—Total number and rate of multilink fragments received and transmitted.</li> <li>• <b>Non-fragments: Input and Output</b>—Total number and rate of nonfragmented multilink frames received and transmitted.</li> <li>• <b>LFI: Input and Output</b>—Total number and rate of link fragmented and interleaved frames and bytes.</li> </ul> </li> </ul>	<b>extensive</b>
<b>Protocol</b>	Protocol family configured on the logical interface.	<b>detail extensive none</b>
<b>MTU</b>	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> .	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route Table</b>	Routing table in which this address exists. For example, <b>Route table:0</b> refers to inet.0.	<b>detail extensive</b>
<b>Addresses, Flags</b>	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> .	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address on the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support.	<b>detail extensive</b>

## 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      0xffff
  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 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 (Redundant Adaptive Services)

<b>Syntax</b>	<pre>show interfaces <i>rspnumber</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(M Series and T Series routers only) Display status information about the specified redundant adaptive services configuration.
<b>Options</b>	<p><b><i>rspnumber</i></b>—Display standard status information about the specified redundant adaptive services configuration.</p> <p><b><i>brief   detail   extensive   terse</i></b>—(Optional) Display the specified level of output.</p> <p><b><i>descriptions</i></b>—(Optional) Display interface description strings.</p> <p><b><i>media</i></b>—(Optional) Display media-specific information about network interfaces.</p> <p><b><i>snmp-index snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b><i>statistics</i></b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show interfaces extensive (Redundant Adaptive Services) on page 1061</a>
<b>Output Fields</b>	See the output field table for the <a href="#">show interfaces (Adaptive Services)</a> 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)

<b>Syntax</b>	<pre>show interfaces rlsqnumber &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;queue&gt; &lt;routing&gt; &lt;snmp-index snmp-index&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced in Junos OS Release 7.6.
<b>Description</b>	(M Series and T Series routers only) Display status information about the specified redundant link services intelligent queuing (IQ) configuration.
<b>Options</b>	<p><b>rlsqnumber</b>—Redundant link services IQ interface name. The logical interface number range of values is 0 through 127.</p> <p><b>none</b>—Display standard status information about the specified redundant link services IQ configuration.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>queue</b>—(Optional) Display queue information about network interfaces.</p> <p><b>routing</b>—(Optional) Display routing information about network interfaces.</p> <p><b>snmp-index snmp-index</b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show interfaces (Redundant Link Services IQ) on page 1074</a></p> <p><a href="#">show interfaces brief (Redundant Link Services IQ) on page 1074</a></p> <p><a href="#">show interfaces detail (Redundant Link Services IQ) on page 1075</a></p> <p><a href="#">show interfaces extensive (Redundant Link Services IQ) on page 1076</a></p>
<b>Output Fields</b>	Table 38 on page 1063 lists the output fields for the <b>show interfaces</b> (redundant link services IQ) command. Output fields are listed in the approximate order in which they appear.

**Table 38: show interfaces (Redundant Link Services IQ) Output Fields**

Field Name	Field Description	Level of Output
Physical Interface		

Table 38: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Physical interface's index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support.	<b>detail extensive</b>
<b>Link-level type</b>	Encapsulation being used on the physical interface: <b>Multilink-Frame-Relay-UNI-NNI</b> (default), <b>LinkService</b> , <b>Frame-relay</b> , <b>Frame-relay-ccc</b> , or <b>Frame-relay-tcc</b> .	All levels
<b>MTU</b>	Maximum transmission unit size on the physical interface.	All levels
<b>Device flags</b>	Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Interface flags</b>	Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .	<b>detail extensive none</b>
<b>Input rate</b>	(Redundant LSQ) Rate of bits and packets received on the interface.	None specified
<b>Output rate</b>	(Redundant LSQ) Rate of bits and packets transmitted on the interface.	None specified
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>
<b>Traffic statistics</b>	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.	<b>detail extensive</b>

Table 38: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Frame exceptions</b>	<p>Information about framing exceptions. Includes events recorded under <b>Exception Events</b> for each logical interface.</p> <ul style="list-style-type: none"> <li>• <b>Oversized frames</b>—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb (kilobits).</li> <li>• <b>Errored input frames</b>—Number of input frame errors.</li> <li>• <b>Input on disabled link/bundle</b>—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.</li> <li>• <b>Output for disabled link/bundle</b>—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.</li> <li>• <b>Queuing drops</b>—Total number of packets dropped before traffic enters the link services IQ interface. Indicates that the interface is becoming oversubscribed.</li> </ul>	<b>extensive</b>
<b>Buffering exceptions</b>	<p>Information about buffering exceptions. Includes events recorded under <b>Exception Events</b> for each logical interface:</p> <ul style="list-style-type: none"> <li>• <b>Packet data buffer overflow</b>—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity.</li> <li>• <b>Fragment data buffer overflow</b>—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.</li> </ul>	<b>extensive</b>

Table 38: show interfaces (Redundant Link Services IQ) Output Fields (continued)

Field Name	Field Description	Level of Output
<b>Assembly exceptions</b>	<p>(Multilink Frame Relay end-to-end only) Information about assembly exceptions. Includes events recorded under <b>Exception Events</b> 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"> <li>• <b>Fragment timeout</b>—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.</li> <li>• <b>Missing sequence number</b>—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.</li> <li>• <b>Out-of-order sequence number</b>—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.</li> <li>• <b>Out-of-range sequence number</b>—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.</li> </ul>	<b>extensive</b>
<b>Hardware errors (sticky)</b>	<p>(Multilink Frame Relay end-to-end only) Information about hardware errors:</p> <ul style="list-style-type: none"> <li>• <b>Data memory error</b>—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support.</li> <li>• <b>Control memory error</b>—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support.</li> </ul>	<b>extensive</b>
<b>Egress queues</b>	Total number of egress queues supported on the specified interface.	<b>detail extensive none</b>

Table 38: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Queue counters</b>	Queue number and its associated user-configured forwarding class name. <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>	<b>detail extensive none</b>
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface	All levels
<b>Index</b>	Logical interface index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	Logical interface SNMP interface index number.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support	<b>detail extensive</b>
<b>Flags</b>	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Encapsulation</b>	Encapsulation being used: PPP or Multilink PPP.	All levels
<b>Bandwidth</b>	Speed at which the interface is running.	All levels
<b>Bundle options</b>	(Multilink Frame Relay end-to-end interfaces only) <ul style="list-style-type: none"> <li>• <b>MRRU</b>—Configured size of the maximum received reconstructed unit (MRRU): 1500 through 4500 bytes. The default is 1504 bytes.</li> <li>• <b>Drop timer period</b>—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.</li> <li>• <b>Sequence number format</b>—Short sequence number header format (MLPPP only).</li> <li>• <b>Fragmentation threshold</b>—Configured fragmentation threshold: 64 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation.</li> <li>• <b>Links needed to sustain bundle</b>—Minimum number of links to sustain the bundle: 1 through 8.</li> <li>• <b>Multilink classes</b>—Number of multilink classes negotiated.</li> <li>• <b>Link layer overhead</b>—Percentage of bundle bandwidth to be set aside for link-layer overhead.</li> </ul>	<b>detail extensive none</b>

Table 38: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Bundle status</b> (MLPPP) or <b>Multilink class status</b> (MC-MLPPP)	<p>Information about bundle status:</p> <ul style="list-style-type: none"> <li>• <b>Remote MRRU</b>—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed.</li> <li>• <b>Received sequence number</b>—Sequence number for received packets.</li> <li>• <b>Transmitted sequence number</b>—Sequence number for transmitted packets.</li> <li>• <b>Packet drops</b>—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.</li> <li>• <b>Fragment drops</b>—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.</li> <li>• <b>MRRU exceeded</b>—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release.</li> <li>• <b>Fragment timeout</b>—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.</li> <li>• <b>Missing sequence number</b>—A gap was detected in the sequence numbers of fragments on a bundle.</li> <li>• <b>Out-of-order sequence number</b>—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.</li> <li>• <b>Out-of-range sequence number</b>—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.</li> <li>• <b>Packet data buffer overflow</b>—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity.</li> <li>• <b>Fragment data buffer overflow</b>—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.</li> </ul>	detail extensive none

Table 38: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Statistics</b>	<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"> <li>• <b>Bundle</b>—Information for each active bundle link. <ul style="list-style-type: none"> <li>• <b>Fragments: Input and Output</b>—Total number and rate of fragments received and transmitted.</li> <li>• <b>Packets: Input and Output</b>—Total number and rate of packets received and transmitted.</li> <li>• <b>Multilink class</b>—(MC-MLPPP only) Information about multiclass links used in the multilink operation.</li> </ul> </li> <li>• <b>Link</b>—Information about links used in the multilink operation. <ul style="list-style-type: none"> <li>• <b>Link name</b>—Interface name of the link services IQ channel and state information (physical link <b>up</b> or <b>down</b>).</li> <li>• <b>Input and Output</b>—Total number and rate of fragments and packets received and transmitted.</li> </ul> </li> </ul>	<b>detail extensive</b>
<b>NCP state</b>	<p>(PPP) Network Control Protocol state.</p> <ul style="list-style-type: none"> <li>• <b>Conf-ack-received</b>—Acknowledgement was received.</li> <li>• <b>Conf-ack-sent</b>—Acknowledgement was sent.</li> <li>• <b>Conf-req-sent</b>—Request was sent.</li> <li>• <b>Down</b>—NCP negotiation is incomplete (not yet completed or has failed).</li> <li>• <b>Not-configured</b>—NCP is not configured on the interface.</li> <li>• <b>Opened</b>—NCP negotiation is successful.</li> </ul>	<b>detail extensive none</b>
<b>Protocol</b>	Protocol family configured on the logical interface.	<b>detail extensive none</b>
<b>MTU</b>	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 <b>Adjusted</b> .	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route Table</b>	Routing table in which this address exists. For example, <b>Route table:0</b> refers to inet.0.	<b>detail extensive</b>
<b>Flags</b>	Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Addresses, Flags</b>	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> .	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>

Table 38: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address on the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support.	<b>detail extensive</b>
<b>MLPPP Bundle Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Logical interface index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	Logical interface SNMP interface index number.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support.	<b>detail extensive</b>
<b>Flags</b>	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>SNMP-Traps</b>	SNMP trap notifications are enabled.	All levels
<b>Encapsulation</b>	Encapsulation being used: PPP, Multilink PPP or Multilink-FR.	All levels
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .	<b>detail extensive none</b>
<b>Bandwidth</b>	Speed at which the interface is running.	All levels
<b>Bundle links information</b>	Information about the bundled links. <ul style="list-style-type: none"> <li>• <b>Active bundle links</b>—Number of active links.</li> <li>• <b>Removed bundle links</b>—Information about links used in the multilink operation.</li> <li>• <b>Disabled bundle links</b>—Number of disabled links.</li> </ul>	<b>detail extensive none</b>



Table 38: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Bundle options</b>	<p>(Multilink Frame Relay end-to-end interfaces only)</p> <ul style="list-style-type: none"> <li>• <b>MRRU</b>—Configured size of the maximum received reconstructed unit (MRRU): <b>1500</b> through <b>4500</b> bytes. The default is <b>1504</b> bytes.</li> <li>• <b>Drop timer period</b>—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: <b>0</b> through <b>2000</b> milliseconds. Values under 5 ms are not recommended. The default setting is <b>0</b>, which disables the timer.</li> <li>• <b>Inner PPP Protocol field compression</b>—Inner PPP protocol compression is enabled or disabled.</li> <li>• <b>Sequence number format</b>—Short sequence number header format (MLPPP only).</li> <li>• <b>Fragmentation threshold</b>—Configured fragmentation threshold: <b>64</b> through <b>16,320</b> bytes, in integer multiples of <b>64</b> bytes. The default setting is <b>0</b>, which disables fragmentation.</li> <li>• <b>Links needed to sustain bundle</b>—Minimum number of links to sustain the bundle: 1 through <b>8</b>.</li> <li>• <b>Multilink classes</b>—Number of multilink classes negotiated.</li> <li>• <b>Link layer overhead</b>—Percentage of bundle bandwidth to be set aside for link-layer overhead.</li> </ul>	<b>detail extensive none</b>

Table 38: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Bundle status</b> (MLPPP)	<p>Information about bundle status:</p> <ul style="list-style-type: none"> <li>• <b>Received sequence number</b>—Sequence number for received packets.</li> <li>• <b>Transmit sequence number</b>—Sequence number for transmitted packets.</li> <li>• <b>Packet drops</b>—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.</li> <li>• <b>Fragment drops</b>—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.</li> <li>• <b>MRRU exceeded</b>—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release.</li> <li>• <b>Fragment timeout</b>—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.</li> <li>• <b>Missing sequence number</b>—A gap was detected in the sequence numbers of fragments on a bundle.</li> <li>• <b>Out-of-order sequence number</b>—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.</li> <li>• <b>Out-of-range sequence number</b>—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.</li> <li>• <b>Packet data buffer overflow</b>—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity.</li> <li>• <b>Fragment data buffer overflow</b>—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.</li> </ul>	detail extensive none

Table 38: show interfaces (Redundant Link Services IQ) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Statistics</b>	<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"> <li>• <b>Bundle</b>—Information for each active bundle link. <ul style="list-style-type: none"> <li>• <b>Multilink: Input and Output</b>—Total number and rate of multilink frames, bytes, and bits per second received and transmitted.</li> <li>• <b>Network: Input and Output</b>—Total number of multilink frames, bytes, and bits per second received and transmitted.</li> </ul> </li> <li>• <b>Link</b>—Information about links used in the multilink operation. <ul style="list-style-type: none"> <li>• <b>Link name</b> is the interface name of the link services IQ channel and state information (physical link <b>up</b> or <b>down</b>) and up time.</li> <li>• <b>Input and Output</b>—Total number and rate of frames, bytes, and bits per second received and transmitted.</li> </ul> </li> </ul>	<b>extensive</b>
<b>Multilink detail statistics</b>	<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"> <li>• <b>Bundle</b>—Information for the bundle link. <ul style="list-style-type: none"> <li>• <b>Fragments: Input and Output</b>—Total number and rate of multilink fragments received and transmitted.</li> <li>• <b>Non-fragments: Input and Output</b>—Total number and rate of nonfragmented multilink frames received and transmitted.</li> <li>• <b>LFI: Input and Output</b>—Total number and rate of link fragmented and interleaved frames and bytes.</li> </ul> </li> </ul>	<b>extensive</b>
<b>Protocol</b>	Protocol family configured on the logical interface.	<b>detail extensive none</b>
<b>MTU</b>	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 <b>Adjusted</b> .	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route Table</b>	Routing table in which this address exists. For example, <b>Route table:0</b> refers to inet.0.	<b>detail extensive</b>
<b>Addresses, Flags</b>	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> .	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address on the logical interface.	<b>detail extensive none</b>

Table 38: 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 1075](#).

## show interfaces load-balancing

<b>Syntax</b>	show interfaces load-balancing <detail>
<b>Release Information</b>	Command introduced in Junos OS Release 11.4.
<b>Description</b>	Display status information about load balancing on aggregated Multiservices (AMS) interfaces.
<b>Options</b>	<b>none</b> —Display standard information about status of all AMS interfaces. <b>detail</b> —(Optional) Display detailed status of all AMS interfaces.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Aggregated Multiservices Interfaces on page 643</a></li> <li>• <a href="#">Example: Configuring an Aggregated Multiservices Interface (AMS) on page 652</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show interfaces load-balancing on page 1079</a> <a href="#">show interfaces load-balancing detail on page 1079</a>
<b>Output Fields</b>	Table 39 on page 1077 lists the output fields for the <b>show interfaces load-balancing</b> (aggregated Multiservices interfaces) command. Output fields are listed in the approximate order in which they appear.

**Table 39: Aggregated Multiservices show interfaces load-balancing Output Fields**

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the aggregated Multiservices (AMS) interface.	All levels
<b>State</b>	Status of AMS interfaces: <ul style="list-style-type: none"> <li>• <b>Up</b>—Interface is configured and operational.</li> <li>• <b>Coming Up</b>—Interface is becoming operational.</li> <li>• <b>Wait Timer</b>—Interface is waiting for member interfaces (mams) to come online.</li> <li>• <b>Members Seen</b>—Member interfaces (mams) are available.</li> <li>• <b>Wait for Members</b>—Member interfaces (mams) are not available.</li> </ul>	All levels
<b>Last change</b>	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 any member.	All levels
<b>Member count</b>	Number of member PICs (mams) that are part of the aggregated interface.	All levels
<b>Members interface</b>	List of all member PICs (mams) that are part of the aggregated interface.	<b>detail</b>

Table 39: Aggregated Multiservices show interfaces load-balancing Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Weight</b>	Weight associated with each member PIC for load balancing. The minimum weight is 1, maximum weight is 100; default weight is 10.	<b>detail</b>
<b>State</b>	Status of each member PIC (mams) : <ul style="list-style-type: none"><li>• <b>Invalid</b>—Configured interface is not valid.</li><li>• <b>Down</b>—Interface is not operational.</li><li>• <b>Active</b>—Interface is configured and operational.</li><li>• <b>Discard</b>—Interface has been discarded.</li><li>• <b>Inactive</b>—Configured interface is not online.</li><li>• <b>Backup</b>—Interface has been configured as backup.</li></ul>	<b>detail</b>



## 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


<b>Syntax</b>	show interfaces redundancy <brief   detail>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. <b>detail</b> option added in Junos OS Release 10.0.
<b>Description</b>	(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><b>NOTE:</b> When you run the <code>show interfaces redundancy</code> command on an MX80 router, it displays the error message, <code>error:the redundancy-interface-process subsystem is not running</code>. This is because an MX80 router does not have a redundant FPC and does not support link protection.</p> </div>	
<b>Options</b>	<b>brief   detail</b> —(Optional) Display the specified level of output.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show interfaces redundancy on page 1081</a> <a href="#">show interfaces redundancy (Aggregated Ethernet) on page 1081</a> <a href="#">show interfaces redundancy detail on page 1081</a>
<b>Output Fields</b>	Table 40 on page 1080 lists the output fields for the <code>show interfaces redundancy</code> command. Output fields are listed in the approximate order in which they appear.

Table 40: show interfaces redundancy Output Fields

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the redundant adaptive services, link services IQ interfaces, or aggregated Ethernet interfaces.	All levels
<b>State</b>	State of the redundant interface: <b>Not present</b> , <b>On primary</b> , <b>On secondary</b> or <b>Waiting for primary MS PIC</b> .	All levels
<b>Last Change</b>	<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"> <li>• GRES is not configured on the router.</li> <li>• The <b>rlsq</b> interface is configured without the <b>hot-standby</b> or <b>warm-standby</b> statements and the backup <b>lsq</b> interface was active before the switchover.</li> <li>• No logical interfaces are configured or all of the configured logical interfaces are down at the time of the switchover.</li> </ul>	All levels

Table 40: show interfaces redundancy Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Primary</b>	Name of the interface configured to be the primary interface.	All levels
<b>Secondary</b>	Name of the interface configured to be the backup interface.	All levels
<b>Current Status</b>	Physical status of the primary and secondary interfaces.	All levels
<b>Mode</b>	Standby mode.	<b>detail</b>

## 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    1sq-0/3/0  1sq-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    1sq-4/0/0  1sq-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        : 1sq-0/2/0
Secondary      : 1sq-1/2/0
Current status : both up
Mode           : hot-standby

Interface      : rlsq0:0
State          : On primary
Last change    : 00:45:46
Primary        : 1sq-0/2/0:0
Secondary      : 1sq-1/2/0:0
Current status : both up
Mode           : warm-standby

```

## show security pki ca-certificate

<b>Syntax</b>	show security pki ca-certificate <brief   detail> <ca-profile <i>ca-profile-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	Display information about certificate authority (CA) digital certificates installed in the router.
<b>Options</b>	<p><b>none</b>—(Same as brief) Display information about all CA digital certificates.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>ca-profile <i>ca-profile-name</i></b>—(Optional) Display information about only the specified CA profile.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show security pki ca-certificate on page 1083</a> <a href="#">show security pki ca-certificate detail on page 1084</a>
<b>Output Fields</b>	Table 41 on page 1082 lists the output fields for the <b>show security pki ca-certificate</b> command. Output fields are listed in the approximate order in which they appear.

Table 41: show security pki ca-certificate Output Fields

Field Name	Field Description	Level of Output
<b>Certificate identifier</b>	Name of the digital certificate.	All levels
<b>Certificate version</b>	Revision number of the digital certificate.	<b>detail</b>
<b>Serial number</b>	Unique serial number of the digital certificate.	<b>detail</b>
<b>Issued by</b>	Authority that issued the digital certificate.	<b>none brief</b>
<b>Issued to</b>	Device that was issued the digital certificate.	<b>none brief</b>
<b>Issuer</b>	<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"> <li>• <b>Common name</b>—Name of the authority.</li> <li>• <b>Organization</b>—Organization of origin.</li> <li>• <b>Organizational unit</b>—Department within an organization.</li> <li>• <b>State</b>—State of origin.</li> <li>• <b>Country</b>—Country of origin.</li> </ul>	<b>detail</b>

Table 41: show security pki ca-certificate Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Subject</b>	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: <ul style="list-style-type: none"> <li>• <b>Common name</b>—Name of the requestor.</li> <li>• <b>Organization</b>—Organization of origin.</li> <li>• <b>Organizational unit</b>—Department within an organization.</li> <li>• <b>State</b>—State of origin.</li> <li>• <b>Country</b>—Country of origin.</li> </ul>	<b>detail</b>
<b>Validity</b>	Time period when the digital certificate is valid. Values are: <ul style="list-style-type: none"> <li>• <b>Not before</b>—Start time when the digital certificate becomes valid.</li> <li>• <b>Not after</b>—End time when the digital certificate becomes invalid.</li> </ul>	All levels
<b>Public key algorithm</b>	Encryption algorithm used with the private key, such as <b>rsaEncryption(1024 bits)</b> .	All levels
<b>Signature algorithm</b>	Encryption algorithm that the CA used to sign the digital certificate, such as <b>sha1WithRSAEncryption</b> .	<b>detail</b>
<b>Fingerprint</b>	Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate.	<b>detail</b>
<b>Distribution CRL</b>	Distinguished name information and the URL for the certificate revocation list (CRL) server.	<b>detail</b>
<b>Use for key</b>	Use of the public key, such as <b>Certificate signing</b> , <b>CRL signing</b> , <b>Digital signature</b> , or <b>Key encipherment</b> .	<b>detail</b>

## 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

<b>Syntax</b>	show security pki certificate-request <brief   detail> <certificate-id <i>certificate-id-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	Display information about manually generated local digital certificate requests that are stored in the router.
<b>Options</b>	<p><b>none</b>—(same as brief) Display information about all local digital certificate requests.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>certificate-id <i>certificate-id-name</i></b>—(Optional) Display information about only the specified local digital certificate request</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear security pki certificate-request on page 996</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show security pki certificate-request on page 1087</a> <a href="#">show security pki certificate-request detail on page 1087</a>
<b>Output Fields</b>	<a href="#">Table 42 on page 1086</a> lists the output fields for the <b>show security pki certificate-request</b> command. Output fields are listed in the approximate order in which they appear.

**Table 42: show security pki certificate-request Output Fields**

Field Name	Field Description	Level of Output
<b>Certificate identifier</b>	Name of the digital certificate.	All levels
<b>Certificate version</b>	Revision number of the digital certificate.	<b>detail</b>
<b>Issued to</b>	Device that was issued the digital certificate.	<b>none brief</b>
<b>Subject</b>	<p>Details of the digital certificate holder organized using the distinguished name format. Possible subfields are:</p> <ul style="list-style-type: none"> <li><b>Common name</b>—Name of the authority.</li> <li><b>Organization</b>—Organization of origin.</li> <li><b>Organizational unit</b>—Department within an organization.</li> <li><b>State</b>—State of origin.</li> <li><b>Country</b>—Country of origin.</li> </ul>	<b>detail</b>
<b>Alternate subject</b>	Domain name or IP address of the device related to the digital certificate.	<b>detail</b>



Table 42: 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"> <li>• <b>Not before</b>—Time when the digital certificate becomes valid.</li> <li>• <b>Not after</b>—End time when the digital certificate becomes invalid.</li> </ul>	All levels
Public key algorithm	Encryption algorithm used with the private key, such as <b>rsaEncryption(1024 bits)</b> .	All levels
Public key verification status	Public key verification status: <b>Failed</b> or <b>Passed</b> . The <b>detail</b> 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.	<b>detail</b>
Use for key	Use of the public key, such as <b>Certificate signing</b> , <b>CRL signing</b> , <b>Digital signature</b> , or <b>Key encipherment</b> .	<b>detail</b>

## Sample Output

### show security pki certificate-request

```

user@host> show security pki certificate-request
Certificate identifier: local-microsoft-2
Issued to: router2.example.com
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.example.com
Alternate subject: router3.example.com
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

<b>Syntax</b>	show security pki crt <brief   detail> <ca-profile <i>ca-profile-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 8.1.
<b>Description</b>	Display information about the certificate revocation lists (CRLs) that are stored in the router.
<b>Options</b>	<p><b>none</b>—(same as brief) Display information about all CRLs.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>ca-profile <i>ca-profile-name</i></b>—(Optional) Display CRL information about only the specified CA profile.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear security pki crt on page 997</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show security pki crt on page 1089</a> <a href="#">show security pki crt detail on page 1089</a>
<b>Output Fields</b>	Table 43 on page 1088 shows the output fields for the <b>show security pki crt</b> command. Output fields are listed in the approximate order in which they appear.

**Table 43: show security pki crt Output Fields**

Field Name	Field Description	Level of Output
<b>CA profile</b>	Name of the configured CA profile.	All levels
<b>CRL version</b>	Revision number of the certificate revocation list.	All levels
<b>CRL number</b>	Number of the certificate revocation list	All levels
<b>CRL issuer</b>	Device that was issued the certificate revocation list.	All levels
<b>Issuer</b>	<p>Details of the digital certificate holder organized using the distinguished name format. Possible subfields are:</p> <ul style="list-style-type: none"> <li>• <b>Common name</b>—Name of the authority.</li> <li>• <b>Organization</b>—Organization of origin.</li> <li>• <b>Organizational unit</b>—Department within an organization.</li> <li>• <b>State</b>—State of origin.</li> <li>• <b>Country</b>—Country of origin.</li> </ul>	<b>detail</b>
<b>Effective date</b>	Date and time the certificate revocation list becomes valid.	All levels

Table 43: show security pki crl Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Next update</b>	Date and time the router will download the latest version of the certificate revocation list.	All levels
<b>Revocation List</b>	<p>List of digital certificates that have been revoked before their expiration date. Values are:</p> <ul style="list-style-type: none"> <li>• <b>Serial number</b>—Unique serial number of the digital certificate</li> <li>• <b>Revocation date</b>—Date and time that the digital certificate was revoked.</li> </ul>	<b>detail</b>

## Sample Output

### show security pki crl

```
user@host> 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

```
user@host> 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

<b>Syntax</b>	show security pki local-certificate <brief   detail> <certificate-id <i>certificate-id-name</i> > <system-generated>
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	Display information about the local digital certificates and the corresponding public keys installed in the router.
<b>Options</b>	<p><b>none</b>—(same as brief) Display information about all local digital certificates and corresponding public keys.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>certificate-id <i>certificate-id-name</i></b>—(Optional) Display information about only the specified the local digital certificate and corresponding public keys.</p> <p><b>system-generated</b>—(Optional) Auto-generated self-signed certificate.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear security pki local-certificate on page 999</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show security pki local-certificate on page 1091</a> <a href="#">show security pki local-certificate detail on page 1092</a>
<b>Output Fields</b>	<a href="#">Table 44 on page 1090</a> lists the output fields for the <b>show security pki local-certificate</b> command. Output fields are listed in the approximate order in which they appear.

**Table 44: show security pki local-certificate Output Fields**

Field Name	Field Description	Level of Output
<b>Certificate identifier</b>	Name of the digital certificate.	All levels
<b>Certificate version</b>	Revision number of the digital certificate.	<b>detail</b>
<b>Serial number</b>	Unique serial number of the digital certificate.	<b>detail</b>
<b>Issued by</b>	Authority that issued the digital certificate.	<b>none brief</b>
<b>Issued to</b>	Device that was issued the digital certificate.	<b>none brief</b>

Table 44: show security pki local-certificate Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Issuer</b>	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"> <li>• <b>Common name</b>—Name of the authority.</li> <li>• <b>Organization</b>—Organization of origin.</li> <li>• <b>Organizational unit</b>—Department within an organization.</li> <li>• <b>State</b>—State of origin.</li> <li>• <b>Country</b>—Country of origin.</li> </ul>	<b>detail</b>
<b>Subject</b>	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: <ul style="list-style-type: none"> <li>• <b>Common name</b>—Name of the authority.</li> <li>• <b>Organization</b>—Organization of origin.</li> <li>• <b>Organizational unit</b>—Department within an organization.</li> <li>• <b>State</b>—State of origin.</li> <li>• <b>Country</b>—Country of origin.</li> </ul>	<b>detail</b>
<b>Alternate subject</b>	Domain name or IP address of the device related to the digital certificate.	<b>detail</b>
<b>Validity</b>	Time period when the digital certificate is valid. Values are: <ul style="list-style-type: none"> <li>• <b>Not before</b>—Start time when the digital certificate becomes valid.</li> <li>• <b>Not after</b>—End time when the digital certificate becomes invalid.</li> </ul>	All levels
<b>Public key algorithm</b>	Encryption algorithm used with the private key, such as <b>rsaEncryption (1024 bits)</b> .	All levels
<b>Public key verification status</b>	Public key verification status: <b>Failed</b> or <b>Passed</b> . The <b>detail</b> output also provides the verification hash.	All levels
<b>Signature algorithm</b>	Encryption algorithm that the CA used to sign the digital certificate, such as <b>sha1WithRSAEncryption</b> .	<b>detail</b>
<b>Fingerprint</b>	Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate.	<b>detail</b>
<b>Distribution CRL</b>	Distinguished name information and URL for the certificate revocation list (CRL) server.	<b>detail</b>
<b>Use for key</b>	Use of the public key, such as <b>Certificate signing</b> , <b>CRL signing</b> , <b>Digital signature</b> , or <b>Key encipherment</b> .	<b>detail</b>

## Sample Output

### show security pki local-certificate

```

user@host> show security pki local-certificate
Certificate identifier: local-entrust2
Issued to: router2.example.com, 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.example.com
Alternate subject: router3.example.com
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

<b>Syntax</b>	<pre>show services cos statistics &lt;brief   detail   extensive&gt; &lt;diffserv   forwarding-class&gt; &lt;interface <i>interface-name</i>&gt; &lt;service-set <i>service-set-name</i>&gt; &lt;summary&gt;</pre>
<b>Release Information</b>	Command introduced in Junos OS Release 8.1.
<b>Description</b>	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.
<b>Options</b>	<p><b>none</b>—Display all services CoS statistics.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>diffserv   forwarding-class</b>—(Optional) Display only the selected information, either DiffServ codepoints or forwarding classes.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display statistics for the specified interface only.</p> <p><b>service-set <i>service-set-name</i></b>—(Optional) Display statistics for the specified service set only.</p> <p><b>summary</b>—(Optional) Display summary of statistics on a per-interface basis.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services cos statistics on page 1094</a> <a href="#">show services cos statistics brief on page 1095</a> <a href="#">show services cos statistics detail on page 1095</a> <a href="#">show services cos statistics extensive on page 1095</a>
<b>Output Fields</b>	Table 45 on page 1093 describes the output fields for the <b>show services cos statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 45: show services cos statistics Output Fields**

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of interface.	All levels
<b>Service set</b>	Name of service set.	All levels
<b>DSCP</b>	DiffServ code point bit pattern.	All levels
<b>Packets in</b>	Number of packets received.	All levels

Table 45: show services cos statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Packets out</b>	Number of packets transmitted.	All levels
<b>Forwarding class</b>	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

<b>Syntax</b>	show services crtp <extensive> <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display Compressed Real-Time Transport Protocol (CRTP) extensive output.
<b>Options</b>	<p><b>none</b>—Display CRTP extensive output for all interfaces.</p> <p><b>extensive</b>—(Optional) Display extensive CRTP information.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display CRTP flow statistics for the specified interface. On M Series and T Series routers, a link services IQ (<b>lsq-fpc/pic/port</b>) or redundant link services IQ (<b>rlsq-fpc/pic/port</b>) interface.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services crtp extensive on page 1097</a>
<b>Output Fields</b>	<a href="#">Table 46 on page 1096</a> lists the output fields for the <b>show services crtp</b> command. Output fields are listed in the approximate order in which they appear.

**Table 46: 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 46: show services crtp Output Fields (*continued*)

Field Name	Field Description
<b>Discards</b>	Number of frames that the incoming packet match code discarded because they were not recognized.
<b>Sessions</b>	Total number of active CRTP sessions.
<b>IP bytes</b>	Number of IP bytes sent and received.
<b>Compressed bytes</b>	Number of compressed IP header bytes sent and received.
<b>CRTP packets</b>	Number of CRTP packets sent and received.
<b>CUDP/CNTCP packets</b>	Number of compressed UDP packets and compressed non-TCP packets sent and received.
<b>Full header packets</b>	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.
<b>Context state packet</b>	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.
<b>IP packets</b>	Number of IP packets sent and received.
<b>Compressed packets</b>	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

<b>Syntax</b>	show services crtp flows <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display Compressed Real-Time Transport Protocol (CRTP) flows.
<b>Options</b>	<p><b>none</b>—Display CRTP flows for all interfaces.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display CRTP flows for the specified interface. On M Series and T Series routers, a link services IQ (<b>lsq-<i>fpc/pic/port</i></b>) or redundant link services IQ (<b>rlsq-<i>fpc/pic/port</i></b>) interface.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services crtp flows on page 1098</a>
<b>Output Fields</b>	<a href="#">Table 47 on page 1098</a> lists the output fields for the <b>show services crtp flows</b> command. Output fields are listed in the approximate order in which they appear.

**Table 47: show services crtp flows Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of the physical interface.
<b>Flow</b>	Received or transmitted flow.
<b>Source</b>	IP source address.
<b>Destination</b>	IP destination address.
<b>SSRC ID</b>	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.
<b>Ctx ID</b>	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 hcm statistics

<b>Syntax</b>	<b>show services hcm statistics rule <i>url-rule-name</i></b>
<b>Description</b>	Display information about statistics associated with the HTTP URL manipulation.
<b>Options</b>	<b><i>url-rule-name</i></b> —Name of the URL rule.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services hcm statistics on page 1100</a>
<b>Output Fields</b>	Table 1 lists the output fields for the <b>show services hcm statistics rule</b> command. Output fields are listed in the approximate order in which they appear.

Table 48: show services hcm statistics rule Output Fields

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the interface.	all
<b>Term id</b>	The ID from the Termination table.	all
<b>Hits</b>	Number of hits for each services hcm statistics rule.	all

## Sample Output

### show services hcm statistics

```

user@host> show services hcm statistics rule url-rule-name
Interface: ms-3/0/0
Term id          Hits
1                10
22              100
333             1000
4444            10000
6               18446744073709551615
Interface: ms-3/1/0
Term id          Hits
1                10
22              100
333             1000
4444            10000
6               18446744073709551615

```

## 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 1105](#)  
[show services ids destination-table extensive on page 1105](#)  
[show services ids destination-table extensive order anomalies on page 1105](#)  
[show services ids pair-table extensive on page 1106](#)  
[show services ids pair-table extensive limit on page 1106](#)  
[show services ids source-table extensive on page 1107](#)  
[show services ids source-table extensive limit on page 1107](#)

**Output Fields** [Table 49 on page 1102](#) lists the output fields for the **show services ids** command. Output fields are listed in the approximate order in which they appear.

**Table 49: show services ids Output Fields**

Field Name	Field Description	Output Level
<b>Interface</b>	Name of an adaptive services interface.	All levels
<b>Service set</b>	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
<b>Sorting order</b>	Primary mode to display information: <b>Anomalies</b> , <b>Bytes</b> , <b>Flows</b> , or <b>Packets</b> .	All levels
<b>Source address</b>	Name of the source address.	All levels
<b>Dest address</b>	Name of the destination address.	All levels
<b>Time</b>	Total time the information has been in the table.	All levels
<b>Flags</b>	<b>Flags</b> can be <b>Forced</b> , <b>F</b> (terse output only), <b>SYNcookie</b> , <b>S</b> (terse output only), <b>Forced+SYNcookie</b> , and <b>F+S</b> (terse output only). The <b>SYNcookie</b> flag is visible only in the destination table.	All levels
<b>Application</b>	Configured application, such as <b>FTP</b> or <b>Telnet</b> .	All levels
<b>Bytes</b>	Total number of bytes sent from the source to the destination address, in thousands ( <b>k</b> ) or millions ( <b>m</b> ).	All levels
<b>Packets</b>	Total number of packets sent from the source to the destination address, in thousands ( <b>k</b> ) or millions ( <b>m</b> ).	All levels
<b>Flows</b>	Total number of flows of packets sent from the source to the destination address, in thousands ( <b>k</b> ) or millions ( <b>m</b> ).	All levels



Table 49: show services ids Output Fields (*continued*)

Field Name	Field Description	Output Level
<b>Anomalies</b>	Total number of packets in the anomaly table, in thousands (k) or millions (m).	All levels
<b>Anomaly description</b>	<p>One or more of the following types of anomalies. For more information, see the detailed descriptions in the stateful firewall section of the <a href="#">System Log Explorer</a>.</p> <ul style="list-style-type: none"> <li>• First packet of TCP session not SYN</li> <li>• ICMP echo request dropped, because sequence number duplicated</li> <li>• ICMP echo reply dropped. No matching sequence number</li> <li>• ICMP echo request dropped. Too many echo requests without echo reply</li> <li>• ICMP header length check failed</li> <li>• ICMP packet length greater than 64K</li> <li>• IP fragment assembly timeout</li> <li>• IP fragment length error</li> <li>• IP fragment overlap</li> <li>• IP packet length greater than 64K</li> <li>• IP packet too short</li> <li>• IP packet with broadcast destination address</li> <li>• IP packet with checksum error</li> <li>• IP packet with incorrect length</li> <li>• IP packet with TTL equal to 0</li> </ul>	extensive

Table 49: show services ids Output Fields (*continued*)

Field Name	Field Description	Output Level
Anomaly description (continued)	<ul style="list-style-type: none"> <li>• IP packet with version other than 4</li> <li>• Land attack (IP src address = dest address)</li> <li>• No matching SFW rule; attempting to create discard flow</li> <li>• Number of open sessions exceeds IDS limit; packet dropped</li> <li>• Packet rate exceeds IDS limit; packet dropped</li> <li>• Session creation rate exceeds IDS limit; packet dropped</li> <li>• SFW application message too long</li> <li>• SFW discard packet contains non-configured IP option types</li> <li>• SFW drop packet because of discard flow</li> <li>• SFW dropped TCP watch packet</li> <li>• SFW rules request FTP active mode data packets to be accepted; attempting to create forward flow</li> <li>• SFW rules request FTP passive mode data packets to be accepted; attempting to create forward flow</li> <li>• SFW rules request packet to be accepted; attempting to create forward or watch flow</li> <li>• SFW rules request packet to be discarded; attempting to create discard flow</li> <li>• SFW rules request packet to be rejected; attempting to create reject flow</li> <li>• SFW discard flow requires packet to be dropped</li> <li>• SFW SYN defense</li> <li>• Smurf attack (ping to IP broadcast address)</li> <li>• TCP FIN/RST or SYN/(URG FIN RST) flags set</li> <li>• TCP header length check failed</li> <li>• TCP port scan (port not in LISTEN state)</li> <li>• TCP seq number zero and FIN/PSH/RST flags set</li> <li>• TCP seq number zero and no flags set</li> <li>• TCP source or destination port zero</li> <li>• TCP SYN flood attack</li> <li>• UDP header length check failed</li> <li>• UDP port scan (port not in LISTEN state)</li> <li>• UDP source or destination port zero</li> </ul>	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 49: 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 <b>show ids services</b> 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      160.0 k    0         14s
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: ss1
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      Count  Rate(eps)  Elapsed
creating forward or watch flow      1      15.6      1m28s
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

<b>Syntax</b>	show services inline nat pool <pool <i>pool--name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 11.4.
<b>Description</b>	Display information about inline Network Address Translation (NAT) pool.
<b>Options</b>	<i>pool-name</i> —Display information about the specified services-inline interface NAT pool.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services inline nat pool on page 1109</a> <a href="#">show services inline nat pool (Network Prefix Translation for IPv6) on page 1110</a>
<b>Output Fields</b>	<a href="#">Table 50 on page 1109</a> lists the output fields for the <b>show services inline nat pool</b> command. Output fields are listed in the order in which they appear.

Table 50: show services inline nat pool Output Fields

Field Name	Field Description
<b>Interface</b>	Name of an <b>si</b> interface hosted on a Trio-based line card.
<b>NAT pool</b>	Name of the pool used for address translations.
<b>Translation type</b>	Translation type specified in the applicable NAT rule for the service set.
<b>Address range</b>	Starting and ending public NAT addresses available for translation.
<b>NATed packets</b>	Number of packets translated for the specified pool.
<b>un-NATed packets</b>	Number of received packets that were not translated.
<b>deNATed packets</b>	Number of packets that were not translated for the specified service PIC.
<b>Errors</b>	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 pool (Network Prefix Translation for IPv6)

```
user@host> show services inline nat pool ss_nptv6_pool1
```

```
Interface: si-4/0/0, Service set: ss_nptv6  
  NAT pool: ss_nptv6_pool1, Translation type: NPTV6  
    Address range: abcd:ef12:3456::/48  
    NATed packets: 0, deNATed packets: 0, Errors: 0
```



## show services inline nat statistics

<b>Syntax</b>	show services inline nat statistics <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 11.4.
<b>Description</b>	Display information about inline Network Address Translation (NAT) address translations.
<b>Options</b>	<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.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services inline nat statistics on page 1112</a> <a href="#">show services inline nat statistics (Network Prefix Translation for IPv6) on page 1112</a>
<b>Output Fields</b>	<a href="#">Table 51 on page 1111</a> lists the output fields for the <b>show services inline nat statistics</b> command. Output fields are listed in the order in which they appear.

**Table 51: 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
Service PIC Name	FPC and PIC slots for the service PIC on which NAT processing is performed	All levels
Data Plane Statistics	Information about packets processed by the data plane for NAT operations	All levels
Control Plane Statistics	Information about packets processed by the control plane for NAT operations	All levels
ICMPv4 errors packets pass through	Number of ICMPv4 error packets that were passed through without being subjected to rules	All levels
ICMPv4 errors packets locally generated	Number of ICMPv4 error packets that were locally generated	All levels

Table 51: show services inline nat statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
ICMPv6 errors packets pass through	Number of ICMPv6 error packets that were passed through without being subjected to rules	All levels
ICMPv6 errors packets locally generated	Number of ICMPv6 error packets that were locally generated	All levels
Dropped packets	Number of packets dropped during inline NAT processing	All levels
NATed packets	Number of packets translated for the specified service PIC.	All levels
deNATed packets	Number of packets that were not translated for the specified service PIC.	All levels
Errors	Number of packets with translation errors.	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 inline nat statistics (Network Prefix Translation for IPv6)

```

user@host> show services inline nat statistics
Service PIC Name                               :si-4/0/0

Control Plane Statistics
  ICMPv4 errors packets pass through           :0
  ICMPv4 errors packets locally generated       :0
  ICMPv6 errors packets pass through           :0
  ICMPv6 errors packets locally generated       :0
  Dropped packets                             :0

Data Plane Statistics
  NATed packets                               :0
  deNATed packets                            :0
  Errors                                      :0

Service PIC Name
:si-4/1/0

Control Plane Statistics
  ICMPv4 errors packets pass through           :0
  ICMPv4 errors packets locally generated       :0
  ICMPv6 errors packets pass through           :0
  ICMPv6 errors packets locally generated       :0
  Dropped packets                             :0

```

Data Plane Statistics	
NATed packets	:0
deNATed packets	:0
Errors	:0

## show services inline software statistics

<b>Syntax</b>	<code>show services inline software statistics</code> <code>&lt;interface <i>interface-name</i>&gt;</code> <code>&lt;v6rd&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 13.3R3.
<b>Description</b>	Display information about inline software activity. The initial implementation of this command reports only on 6rd activity.
<b>Options</b>	<p><b>interface <i>interface-name</i></b>—(Optional) Display information about the specified services-inline interface only. When a specific interface is not specified, statistics for all services-inline interfaces are shown.</p> <p><b>v6rd</b>—(Optional) Display information for 6rd.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show services inline software statistics on page 1115</a></p> <p><a href="#">show services inline software statistics interface on page 1116</a></p>
<b>Output Fields</b>	Table 52 on page 1114 lists the output fields for the <b>show services inline software statistics</b> command. Output fields are listed in the order in which they appear.

**Table 52: show services inline software statistics Output Fields**

Field Name	Field Description
Service PIC Name	Name of the service PIC for which statistics are displayed.
<b>Control Plane Statistics</b>	
ICMPv4 echo requests to software concentrator	Number of ICMPv4 echo received by the software concentrator.
ICMPv4 echo responses from software concentrator	Number of ICMPv4 echo responses sent from the software concentrator.
Dropped ICMPv4 packets to software concentrator	Number of ICMP packets (except ICMP request) received by the software concentrator. All these packets are dropped in by the packet forwarding engine Ukernel.
Trace route UDP packets to software concentrator	Number of UDP trace route packets (port numbers 33434 through 33534) received by the software concentrator.

Table 52: show services inline software statistics Output Fields (*continued*)

Field Name	Field Description
ICMPv4 Port unreachable errors sent from software concentrator	Number of ICMP port unreachable errors sent by the software concentrator after receiving the UDP trace route packets.
Other dropped IPv4 packets to software concentrator	Number of non-ICMP packets that were received and dropped because of fragmentation during encapsulation or decapsulation.
<b>Data Plane Statistics</b>	
6rd decaps	Number of 6rd decapsulated packets and bytes in the data plane. Decapsulation includes removing the outer IPv4 header and routing the inner IPv6 packet.
6rd encaps	Number of 6rd encapsulated (IPv4) packets and bytes in the data plane.
6rd decap errors	Number of all the packets and bytes that are not IPv4-IPv6, IPv4-UDP, or IPv4-ICMP packets.
6rd decap fragment errors	Number of IPv4 fragmented packets and bytes.
6rd decap spoof attacks	Number of spoof attack packets and bytes, which includes packets for which the 6rd derived IPv4 address does not match with the source IPv4 address and packets for which the source IPv6 prefix does not match the 6rd IPv6 prefix.
6rd encaps v4 mtu errors	Count of packets and bytes with IPv4 encapsulation MTU errors. For downlink packets after encapsulating with an IPv4 header, if the packet length is more than Tunnel MTU then it is dropped as v4 MTU errors. For these packet drops, an ICMPv6 packet too big error is sent back to the sender.

## Sample Output

### show services inline software statistics

```

user@host> show services inline software statistics
Border Router v6rd statistics:

Service PIC Name                               si-0/0/0

Control Plane Statistics
  ICMPv4 echo requests to software concentrator      0
  ICMPv4 echo responses from software concentrator   0
  Dropped ICMPv4 packets to software concentrator    0
  Trace route UDP packets to software concentrator   0
  ICMPv4 Port unreachable errors sent from software concentrator 0
  Other dropped IPv4 packets to software concentrator 0

Data Plane Statistics
  6rd decaps      Packets      Bytes
32222173891      3061106519645

```

6rd encaps	415480622	28252710148
6rd decap errors	0	0
6rd decap fragment errors	0	0
6rd decap spoof attacks	0	0

Service PIC Name	si-0/2/0
------------------	----------

Control Plane Statistics		
ICMPv4 echo requests to software concentrator		0
ICMPv4 echo responses from software concentrator		0
Dropped ICMPv4 packets to software concentrator		0
Trace route UDP packets to software concentrator		0
ICMPv4 Port unreachable errors sent from software concentrator		0
Other dropped IPv4 packets to software concentrator		0

Data Plane Statistics	Packets	Bytes
6rd decaps	0	0
6rd encaps	0	0
6rd decap errors	0	0
6rd decap fragment errors	0	0
6rd decap spoof attacks	0	0
6rd encaps v4 mtu errors	0	0

#### **show services inline software statistics interface**

```
user@host> show services inline software statistics interface si-0/0/0
```

## show services ipsec-vpn certificates

<b>Syntax</b>	show services ipsec-vpn certificates <brief   detail> <service-set <i>service-set</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	(Adaptive services interfaces only) Display local and remote certificates installed in the IPsec configuration memory cache that are used for the IKE negotiation.
<b>Options</b>	<p><b>none</b>—(same as brief) Display information about local and remote certificates associated with all service sets.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>service-set <i>service-set</i></b>—(Optional) Display information about local and remote certificates associated with only the specified service set.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show security ipsec-vpn certificates on page 1118</a> <a href="#">show security ipsec-vpn certificates detail on page 1118</a>
<b>Output Fields</b>	Table 53 on page 1117 lists the output fields for the <b>show services ipsec-vpn certificates</b> command. Output fields are listed in the approximate order in which they appear.

Table 53: show services ipsec-vpn certificates Output Fields

Field Name	Field Description	Level of Output
<b>Service set</b>	Name of the IPsec service set.	All levels
<b>Total entries</b>	Number of certificate cache entries.	All levels
<b>Certificate cache entry</b>	Identification number of the certificate cache entry.	All levels
<b>Flags</b>	Information about the digital certificate, including whether the certificate is a root certificate and trusted.	none <b>brief</b>
<b>Issued to</b>	Device that was issued the digital certificate.	none <b>brief</b>
<b>Issued by</b>	Authority that issued the digital certificate.	none <b>brief</b>
<b>Certificate version</b>	Revision number of the digital certificate.	<b>detail</b>
<b>Serial number</b>	Unique serial number of the digital certificate.	<b>detail</b>
<b>Alternate subject</b>	Domain name or IP address of the device related to the digital certificate.	All levels

Table 53: 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"> <li>• <b>Not before</b>—Start time when the digital certificate becomes valid.</li> <li>• <b>Not after</b>—End time when the digital certificate becomes invalid.</li> </ul>	none <b>brief</b>
Public key algorithm	Specifies the encryption algorithm used with the private key, such as <b>rsaEncryption (1024 bits)</b> .	<b>detail</b>
Signature algorithm	Encryption algorithm that the CA used to sign the digital certificate, such as <b>sha1WithRSAEncryption</b> .	<b>detail</b>
Fingerprint	Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate.	<b>detail</b>
Distribution CRL	Distinguished name information and the URL for the certificate revocation list (CRL) server.	<b>detail</b>
Use for key	Use of the public key, such as <b>Certificate signing</b> , <b>CRL signing</b> , <b>Digital signature</b> , or <b>Key encipherment</b> .	<b>detail</b>

## 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.example.com, Issued by: juniper
  Alternate subject: router3.example.com
  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.example.com, Issued by: juniper
  Alternate subject: router2.example.com
  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.example.com
  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.example.com
  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

<b>Syntax</b>	show services ipsec-vpn ike security-associations <brief   detail> <peer-address>
<b>Release Information</b>	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.
<b>Description</b>	(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.
<b>Options</b>	<b>none</b> —(same as brief) Display standard information for all IPsec security associations.  <b>brief   detail</b> —(Optional) Display the specified level of output.  <b>peer-address</b> —(Optional) Display information about a particular security association address.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services ipsec-vpn ike security-associations on page 1122</a> <a href="#">show services ipsec-vpn ike security-associations detail on page 1123</a>
<b>Output Fields</b>	<a href="#">Table 54 on page 1120</a> lists the output fields for the <b>show services ipsec-vpn ike security-associations</b> command. Output fields are listed in the approximate order in which they appear.

**Table 54: show services ipsec-vpn ike security-associations Output Fields**

Field Name	Field Description	Level of Output
<b>IKE peer</b>	Remote end of the IKE negotiation.	<b>detail</b>
<b>Role</b>	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.	<b>detail</b>
<b>Remote Address</b>	Responder's address.	none specified
<b>State</b>	State of the IKE security association: <ul style="list-style-type: none"> <li>• <b>Matured</b>—IKE security association is established.</li> <li>• <b>Not matured</b>—The IKE security association is in the process of negotiation.</li> </ul>	none specified
<b>Initiator cookie</b>	When the IKE negotiation is triggered, a random number is sent to the remote node.	All levels

Table 54: show services ipsec-vpn ike security-associations Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Responder cookie</b>	<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
<b>Exchange type</b>	<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"> <li>• <b>Main</b>—The exchange is done with six messages. <b>Main</b> encrypts the payload, protecting the identity of the neighbor.</li> <li>• <b>Aggressive</b>—The exchange is done with three messages. <b>Aggressive</b> does not encrypt the payload, leaving the identity of the neighbor unprotected.</li> <li>• <b>IKEv2</b>—The exchange is negotiated using IKE version 2.</li> </ul>	All levels
<b>PIC</b>	The services PIC for which the IKE security associations are displayed.	All levels
<b>Authentication method</b>	Type of authentication determines which payloads are exchanged and when they are exchanged. The Junos OS supports only <b>pre-shared keys</b> .	<b>detail</b>
<b>Local</b>	Prefix and port number of the local end.	<b>detail</b>
<b>Remote</b>	Prefix and port number of the remote end.	<b>detail</b>
<b>Lifetime</b>	Number of seconds remaining until the IKE security association expires.	<b>detail</b>
<b>Algorithms</b>	<p>Header for the IKE algorithms output.</p> <ul style="list-style-type: none"> <li>• <b>Authentication</b>—(<b>detail</b> output only) Type of authentication algorithm used: <b>md5</b> or <b>sha1</b></li> <li>• <b>Encryption</b>—(<b>detail</b> output only) Type of encryption algorithm used: <b>des-cbc</b>, <b>3des-cbc</b>, or <b>None</b>.</li> <li>• <b>Pseudo random function</b>—Function that generates highly unpredictable random numbers: <b>hmac-md5</b> or <b>hmac-sha1</b>.</li> </ul>	<b>detail</b>
<b>Traffic statistics</b>	<p>Number of bytes and packets received and transmitted on the IKE security association.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the IKE security association.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the IKE security association.</li> </ul>	<b>detail</b>

Table 54: show services ipsec-vpn ike security-associations Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Flags</b>	Notification to the key management process of the status of the IKE negotiation: <ul style="list-style-type: none"> <li><b>caller notification sent</b>—Caller program notified about the completion of the IKE negotiation.</li> <li><b>waiting for done</b>—Negotiation is done. The library is waiting for the remote end retransmission timers to expire.</li> <li><b>waiting for remove</b>—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.</li> <li><b>waiting for policy manager</b>—Negotiation is waiting for a response from the policy manager.</li> </ul>	<b>detail</b>
<b>IPsec security associates</b>	Number of IPsec security associations created and deleted with this IKE security association.	<b>detail</b>
<b>Phase 2 negotiations in progress</b>	Number of phase 2 negotiations in progress and status information: <ul style="list-style-type: none"> <li>Negotiation type—Type of phase 2 negotiation. The Junos OS currently supports <b>quick mode</b>.</li> <li>Message ID—Unique identifier for a phase 2 negotiation.</li> <li>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>.</li> <li>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>.</li> <li>Flags—Notification to the key management process of the status of the IKE negotiation: <ul style="list-style-type: none"> <li><b>caller notification sent</b>—Caller program notified about the completion of the IKE negotiation.</li> <li><b>waiting for done</b>—Negotiation is done. The library is waiting for the remote end retransmission timers to expire.</li> <li><b>waiting for remove</b>—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.</li> <li><b>waiting for policy manager</b>—Negotiation is waiting for a response from the policy manager.</li> </ul> </li> </ul>	<b>detail</b>

## 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

<b>Syntax</b>	show services ipsec-vpn ipsec security-associations <brief   detail   extensive> <service-set <i>service-set-name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(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.
<b>Options</b>	<p><b>none</b>—Display standard information about IPsec security associations for all service sets.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>service-set <i>service-set-name</i></b>—(Optional) Display information about a particular service set.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services ipsec-vpn ipsec security associations extensive on page 1127</a>
<b>Output Fields</b>	<a href="#">Table 55 on page 1124</a> lists the output fields for the <b>show services ipsec-vpn ipsec security-associations</b> command. Output fields are listed in the approximate order in which they appear.

**Table 55: show services ipsec-vpn ipsec security-associations Output Fields**

Field Name	Field Description	Level of Output
<b>Service set</b>	Name of the service set for which the IPsec security associations are defined. If appropriate, includes the outside service interface VRF name.	All levels
<b>Rule</b>	Name of the rule set applied to the security association.	<b>detail extensive</b>
<b>Term</b>	Name of the IPsec term applied to the security association.	<b>detail extensive</b>
<b>Tunnel index</b>	Numeric identifier of the specific IPsec tunnel for the security association.	<b>detail extensive</b>
<b>Local gateway</b>	Gateway address of the local system.	All levels
<b>Remote gateway</b>	Gateway address of the remote system.	All levels
<b>IPsec inside interface</b>	Name of the logical interface hosting the IPsec tunnels.	All levels
<b>Tunnel MTU</b>	MTU of the IPsec tunnel.	All levels

Table 55: show services ipsec-vpn ipsec security-associations Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Local identity</b>	<p>Protocol, address or prefix, and port number of the local entity of the IPsec association. The format is <b>id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)</b>. 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 <b>id-data-len</b> parameter can be one of the following, depending on the address configured in the IPsec rule:</p> <ul style="list-style-type: none"> <li>For an IPv4 address, the length is 4 and the value displayed is 3.</li> <li>For a subnet mask of an IPv4 address, the length is 8 and the value displayed is 7.</li> <li>For a range of IPv4 addresses, the length is 8 and the value displayed is 7.</li> <li>For an IPv6 address prefix, the length is 16 and the value displayed is 15.</li> <li>For a subnet mask of an IPv6 address prefix, the length is 32 and the value displayed is 31.</li> <li>For a range of IPv6 address prefixes, the length is 32 and the value displayed is 31.</li> </ul> <p>The value of the <b>id-data-presentation</b> 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
<b>Remote identity</b>	<p>Protocol, address or prefix, and port number of the remote entity of the IPsec association. The format is <b>id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)</b>. 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 <b>id-data-len</b> parameter can be one of the following, depending on the address configured in the IPsec rule:</p> <ul style="list-style-type: none"> <li>For an IPv4 address, the length is 4 and the value displayed is 3.</li> <li>For a subnet mask of an IPv4 address, the length is 8 and the value displayed is 7.</li> <li>For a range of IPv4 addresses, the length is 8 and the value displayed is 7.</li> <li>For an IPv6 address prefix, the length is 16 and the value displayed is 15.</li> <li>For a subnet mask of an IPv6 address prefix, the length is 32 and the value displayed is 31.</li> <li>For a range of IPv6 address prefixes, the length is 32 and the value displayed is 31.</li> </ul> <p>The value of the <b>id-data-presentation</b> 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
<b>Primary remote gateway</b>	IP address of the configured primary remote peer.	All levels
<b>Backup remote gateway</b>	IP address of the configured backup remote peer.	All levels

Table 55: show services ipsec-vpn ipsec security-associations Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	State of the primary or backup interface: <b>Active</b> , <b>Offline</b> , or <b>Standby</b> . Both ES PICs are initialized to <b>Offline</b> . For primary and backup peers, <b>State</b> can be <b>Active</b> or <b>Standby</b> . If both peers are in a state of <b>Standby</b> , no connection exists yet between the two peers.	All levels
<b>Failover counter</b>	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
<b>Direction</b>	Direction of the security association: <b>inbound</b> or <b>outbound</b> .	All levels
<b>SPI</b>	Value of the security parameter index.	All levels
<b>AUX-SPI</b>	Value of the auxiliary security parameter index. <ul style="list-style-type: none"> <li>When the value of <b>Protocol</b> is <b>AH</b> or <b>ESP</b>, <b>AUX-SPI</b> is always 0.</li> <li>When the value of <b>Protocol</b> is <b>AH+ESP</b>, <b>AUX-SPI</b> is always a positive integer.</li> </ul>	All levels
<b>Mode</b>	Mode of the security association: <ul style="list-style-type: none"> <li><b>transport</b>—Protects single host-to-host protections.</li> <li><b>tunnel</b>—Protects connections between security gateways.</li> </ul>	<b>detail extensive</b>
<b>Type</b>	Type of security association: <ul style="list-style-type: none"> <li><b>manual</b>—Security parameters require no negotiation. They are static, and are configured by the user.</li> <li><b>dynamic</b>—Security parameters are negotiated by the IKE protocol. Dynamic security associations are not supported in transport mode.</li> </ul>	<b>detail extensive</b>
<b>State</b>	Status of the security association: <ul style="list-style-type: none"> <li><b>Installed</b>—The security association is installed in the security association database. (For transport mode security associations, the value of State must always be Installed.)</li> <li><b>Not installed</b>—The security association is not installed in the security association database.</li> </ul>	<b>detail extensive</b>
<b>Protocol</b>	Protocol supported: <ul style="list-style-type: none"> <li><b>transport</b> mode supports Encapsulation Security Protocol (<b>ESP</b>) or Authentication Header (<b>AH</b>).</li> <li><b>tunnel</b> mode supports <b>ESP</b> or <b>AH+ESP</b>.</li> </ul>	All levels
<b>Authentication</b>	Type of authentication used: <b>hmac-md5-96</b> , <b>hmac-sha1-96</b> , or <b>none</b> .	<b>detail extensive</b>
<b>Encryption</b>	Type of encryption algorithm used: can be <b>aes-cbc (128 bits)</b> , <b>aes-cbc (192 bits)</b> , <b>aes-cbc (256 bits)</b> , <b>des-cbc</b> , <b>3des-cbc</b> , or <b>None</b> .	<b>detail</b>



Table 55: show services ipsec-vpn ipsec security-associations Output Fields (*continued*)

Field Name	Field Description	Level of Output
Soft lifetime Hard lifetime	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. <ul style="list-style-type: none"> <li>Expires in <i>seconds</i> <b>seconds</b>—Number of seconds left until the security association expires.</li> <li>Expires in <i>kilobytes</i> <b>kilobytes</b>—Number of kilobytes left until the security association expires.</li> </ul>	detail extensive
Anti-replay service	State of the service that prevents packets from being replayed: <b>Enabled</b> or <b>Disabled</b> .	detail extensive
Replay window size	Configured size, in packets, of the antireplay service window: <b>32</b> or <b>64</b> . The antireplay window size protects the receiver against replay attacks by rejecting old or duplicate packets. If the replay window size is <b>0</b> , 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

<b>Syntax</b>	show services ipsec-vpn ipsec statistics <brief   detail> <remote-gw remote-peer-address> <service-set service-set-name>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. New fields added in Junos OS Release 10.0.
<b>Description</b>	(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.
<b>Options</b>	<p><b>none</b>—Display standard IPsec statistics for all service sets.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>remote-gw remote-peer-address</b>—(Optional) Display IPsec statistics for an individual IPsec tunnel and an individual remote host.</p> <p><b>service-set service-set-name</b>—(Optional) Display information about a particular service set.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services ipsec-vpn ipsec statistics detail on page 1130</a> <a href="#">show services ipsec-vpn ipsec statistics remote-gw on page 1130</a>
<b>Output Fields</b>	<a href="#">Table 56 on page 1128</a> lists the output fields for the <b>show services ipsec-vpn ipsec statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 56: show services ipsec-vpn ipsec statistics Output Fields**

Field Name	Field Description	Level of Output
<b>PIC</b>	The physical interface on which the IPsec tunnel is configured.	All levels
<b>Service set</b>	Name of the service set for which the IPsec tunnel is defined.	All levels
<b>Local gateway</b>	Gateway address of the local system.	All levels
<b>Remote gateway</b>	Gateway address of the remote system.	All levels
<b>Tunnel index</b>	Numeric identifier of the specific IPsec tunnel for the security association.	All levels

Table 56: show services ipsec-vpn ipsec statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>ESP statistics</b>	Encapsulation Security Payload (ESP) statistics: <ul style="list-style-type: none"> <li>• <b>Encrypted bytes</b>—Total number of bytes encrypted by the local system across the IPsec tunnel.</li> <li>• <b>Decrypted bytes</b>—Total number of bytes decrypted by the local system across the IPsec tunnel.</li> <li>• <b>Encrypted packets</b>—Total number of packets encrypted by the local system across the IPsec tunnel.</li> <li>• <b>Decrypted packets</b>—Total number of packets decrypted by the local system across the IPsec tunnel.</li> </ul>	All levels
<b>AH Statistics</b>	Authentication Header statistics: <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Total number of bytes received by the local system across the IPsec tunnel.</li> <li>• <b>Output bytes</b>—Total number of bytes transmitted by the local system across the IPsec tunnel.</li> <li>• <b>Input packets</b>—Total number of packets received by the local system across the IPsec tunnel.</li> <li>• <b>Output packets</b>—Total number of packets transmitted by the local system across the IPsec tunnel.</li> </ul>	All levels
<b>Errors</b>	<ul style="list-style-type: none"> <li>• <b>AH authentication failures</b>—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.</li> <li>• <b>ESP authentication failures</b>—Number of Encapsulation Security Payload (ESP) failures. An ESP failure occurs when there is an authentication mismatch in ESP packets.</li> <li>• <b>ESP Decryption failures</b>—Number of ESP decryption failures.</li> <li>• <b>Bad headers</b>—Number of invalid headers detected.</li> <li>• <b>Bad trailers</b>—Number of invalid trailers detected.</li> <li>• <b>Replay before window drops</b>—Number of replay errors. A replay error is generated when a duplicate packet is received within the replay window.</li> <li>• <b>Replayed pkts</b>—Number of packets replayed.</li> <li>• <b>IP integrity errors</b>—Number of IP integrity errors.</li> <li>• <b>Exceeds tunnel MTU</b>—Number of times the tunnel maximum transmission unit (MTU) value was exceeded.</li> <li>• <b>Rule lookup failures</b>—Number of rule lookup failures.</li> <li>• <b>No SA errors</b>—Number of errors resulting from a missing security association (SA).</li> <li>• <b>Flow errors</b>—Number of flow errors.</li> <li>• <b>Misc errors</b>—Number of miscellaneous errors.</li> </ul>	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

<b>Syntax</b>	show services link-services cpu-usage <brief   detail> <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 8.4.
<b>Description</b>	(M Series and T Series routers only) Display information about Link Services IQ (LSQ) CPU usage.
<b>Options</b>	<p><b>none</b>—Display standard information about CPU usage for all LSQ interfaces.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display information about the specified LSQ interface.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services link-services cpu-usage brief (AS PIC) on page 1133</a> <a href="#">show services link-services cpu-usage brief (MultiServices PIC) on page 1133</a> <a href="#">show services link-services cpu-usage detail (AS PIC) on page 1133</a> <a href="#">show services link-services cpu-usage detail (MultiServices PIC) on page 1134</a>
<b>Output Fields</b>	Table 57 on page 1131 lists the output fields for the <b>show services link-services cpu-usage</b> command. Output fields are listed in the approximate order in which they appear.

**Table 57: show services link-services cpu-usage Output Fields**

Field Name	Field Description	Level of Output
<b>Role</b>	CPU functional category.	<b>brief</b>
<b>1 Second Average</b>	Percentage of usage during 1-second duration.	All levels
<b>5 Second Average</b>	Percentage of usage during 5-second duration.	All levels
<b>QoS</b>	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
<b>Sequencer</b>	Assigns sequence numbers to outgoing MLPPP fragments and interleaves link fragmentation and interleaving (LFI) traffic.	All levels
<b>Load Balancer</b>	Distributes load across different fragmenter CPUs.	All levels
<b>Fragmenter</b>	Main LSQ CPU; fragments IP packets into MLPPP fragments and also reassembles MLPPP fragments into IP packets.	All levels
<b>Total</b>	Sum of all CPU functions.	<b>brief</b>

Table 57: show services link-services cpu-usage Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Idle</b>	Counts idle cycles when the CPU does not have any work.	<b>detail</b>
<b>Timer</b>	Takes care of periodic events driven by a timer, such as timeouts.	<b>detail</b>
<b>System</b>	System housekeeping thread.	<b>detail</b>
<b>Input (QoS)</b>	Acquires and queues incoming IP frames from hardware interfaces.	<b>detail</b>
<b>Output (QoS)</b>	Sends scheduled frames to the next processing CPU.	<b>detail</b>
<b>Output Frags (QoS)</b>	Sends outstanding frames to the fragmenter CPU.	<b>detail</b>
<b>Bypass (QoS)</b>	Sends outstanding frames for LFI.	<b>detail</b>
<b>Free frame (QoS)</b>	Frees dropped frames.	<b>detail</b>
<b>CPUnumber</b>	Identifier number of specific CPU.	<b>detail</b>
<b>Drop (Fragmenter)</b>	Drops frames that have been marked by the QoS CPU.	<b>detail</b>
<b>Frag (Fragmenter)</b>	Fragments IP frames into MLPPP fragments.	<b>detail</b>
<b>Reass (Fragmenter)</b>	Reassembles MLPPP fragments into IP frames.	<b>detail</b>
<b>Freeback (Fragmenter)</b>	Handles freeback of credits from other CPUs (MultiServices PICs only).	<b>detail</b>
<b>Input LFI (Sequencer)</b>	Receives LFI traffic from QoS CPU and transmits it with strict priority over MLPPP.	<b>detail</b>
<b>Input Frag (Sequencer)</b>	Receives MLPPP fragments from fragmenter CPUs, assigns sequence numbers, and appends MLPPP headers.	<b>detail</b>
<b>Output Frag (Sequencer)</b>	Load-balances and transmits fragments across links.	<b>detail</b>
<b>Retry (Sequencer)</b>	Retries transmission if hardware was busy in the previous attempt.	<b>detail</b>
<b>Input Alloc (Load Balancer)</b>	Acquires frames from hardware interfaces and validates them.	<b>detail</b>
<b>Input (Load Balancer)</b>	Performs error and sanity checks and check frames for PortMapping.	<b>detail</b>
<b>Output (Load Balancer)</b>	Sends frame to next processing CPU.	<b>detail</b>

Table 57: show services link-services cpu-usage Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Freeback</b> (Load Balancer)	Handles freeback of credits from other CPUs.	<b>detail</b>

## 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  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  Input  Output  Free
              Alloc 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

<b>Syntax</b>	show services l2tp multilink <brief   detail   extensive   statistics> <bundle-id <i>number</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(M10i and M7i routers only) Display L2TP output organized by multilink bundle.
<b>Options</b>	<p><b>none</b>—Same as brief.</p> <p><b>brief   detail   extensive   statistics</b>—(Optional) Display the specified level of output. Use the <b>statistics</b> 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><b>bundle-id <i>number</i></b>—(Optional) Display L2TP multilink bundle information for only the specified bundle.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">L2TP Services Configuration Overview on page 676</a></li> <li>• <a href="#">L2TP Minimum Configuration on page 677</a></li> <li>• <a href="#">clear services l2tp multilink on page 978</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show services l2tp multilink extensive on page 1138</a>
<b>Output Fields</b>	<p><a href="#">Table 58 on page 1135</a> lists the output fields for the <b>show services l2tp multilink</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 58: show services l2tp multilink Output Fields**

Field Name	Field Description	Level of Output
<b>Bundle ID</b>	Bundle identifier.	All levels
<b>Links</b>	Number of links in the multilink bundle.	All levels
<b>Bundle endpoint</b>	Endpoint discriminator that represents the device transmitting the packet.	All levels
<b>Input MRRU</b>	Maximum packet size that the input interface can process.	detail
<b>Output MRRU</b>	Maximum packet size that the output interface can process.	detail

Table 58: show services l2tp multilink Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Session local ID</b>	Identifier of the local endpoint of the L2TP session, as assigned by the L2TP network server (LNS).	detail
<b>Session remote ID</b>	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	detail
<b>State</b>	Status of the L2TP session: <ul style="list-style-type: none"> <li>• <b>Established</b>—The session is operating.</li> <li>• <b>closed</b>—The session is being closed.</li> <li>• <b>destroyed</b>—The session is being destroyed.</li> <li>• <b>clean-up</b>—The session is being cleaned up.</li> <li>• <b>lns-ic-accept-new</b>—A new session is being accepted.</li> <li>• <b>lns-ic-idle</b>—The session has been created and is idle.</li> <li>• <b>lns-ic-reject-new</b>—The new session is being rejected.</li> <li>• <b>lns-ic-wait-connect</b>—The session is waiting for the peer's incoming call connected (ICCN) message.</li> </ul>	detail
<b>Username</b>	Name of the user logged in to the session.	detail
<b>Mode</b>	Mode of the interface representing the multilink bundle: <b>dedicated</b> or <b>shared</b> .	extensive
<b>Local IP</b>	IP address of the local endpoint of the Point-to-Point Protocol (PPP) session.	extensive
<b>Remote IP</b>	IP address of the remote endpoint of the PPP session.	extensive
<b>Local name</b>	Name of the LNS instance in which the session was created.	extensive
<b>Remote name</b>	Name of the LAC from which the session was created.	extensive
<b>Local MRU</b>	Maximum receive unit (MRU) setting of the local device, in bytes.	extensive

Table 58: 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"> <li>• <b>Control Tx</b>—Amount of control information transmitted, in packets and bytes.</li> <li>• <b>Control Rx</b>—Amount of control information received, in packets and bytes.</li> <li>• <b>Data Tx</b>—Amount of data transmitted, in packets and bytes.</li> <li>• <b>Data Rx</b>—Amount of data received, in packets and bytes.</li> <li>• <b>Errors Tx</b>—Number of errors transmitted, in packets.</li> <li>• <b>Errors Rx</b>—Number of errors received, in packets.</li> <li>• <b>Lcp Echo Req Tx</b>—Number of LCP echo requests transmitted, in packets.</li> <li>• <b>Lcp Echo Req Rx</b>—Number of LCP echo requests received, in packets.</li> <li>• <b>Lcp Echo Rep Tx</b>—Number of LCP echo responses transmitted, in packets.</li> <li>• <b>Lcp Echo Rep Rx</b>—Number of LCP echo responses received, in packets.</li> <li>• <b>Lcp Echo Req Timeout</b>—Number of LCP echo requests that timed out.</li> <li>• <b>Lcp Echo Req Error</b>—Number of errors received for LCP echo packets.</li> <li>• <b>Lcp Echo Rep Error</b>—Number of errors transmitted for LCP echo packets.</li> <li>• <b>MRRU</b>—Maximum packet size processed.</li> <li>• <b>TX</b>—Number of packets transmitted.</li> <li>• <b>RX</b>—Number of packets received.</li> <li>• <b>link</b>—Link of the multilink bundle associated with the L2TP session.</li> </ul>	extensive

## Sample Output

### show services l2tp multilink extensive

```

user@host> show services l2tp multilink extensive
Bundle ID: 1
  Links: 2, Bundle endpoint: user@example.com
  Input MRRU: 1524, Output MRRU: 1524
  Session local ID: 46122, Session remote ID: 39307
    State: Established, Username: user1@example.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@example.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

<b>Syntax</b>	<pre>show services l2tp radius &lt;accounting (servers   statistics)&gt; &lt;authentication (servers   statistics)&gt; &lt;servers&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0.
<b>Description</b>	(M7i, M10i, and M120 routers only) Display RADIUS servers and statistics information for the RADIUS servers configured on the router.
<b>Options</b>	<p>You must include one of the following keywords to provide a valid completion for the command:</p> <p><b>accounting (servers   statistics)</b>—(Optional) Display RADIUS servers or statistical accounting information only.</p> <p><b>authentication (servers   statistics)</b>—(Optional) Display RADIUS servers or statistical authentication information only.</p> <p><b>servers</b>—(Optional) Display RADIUS authentication and accounting server information only.</p> <p><b>statistics</b>—(Optional) Display RADIUS authentication and accounting statistics information only.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">L2TP Services Configuration Overview on page 676</a></li> <li>• <a href="#">L2TP Minimum Configuration on page 677</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show services l2tp radius servers on page 1141</a> <a href="#">show services l2tp radius statistics on page 1141</a>
<b>Output Fields</b>	<p><a href="#">Table 59 on page 1139</a> lists the output fields for the <b>show services l2tp radius</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 59: show services l2tp radius Output Fields**

Field Name	Field Description
IP Address	IP address of the server.
State	( <b>servers</b> 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	( <b>servers</b> keyword only) Number of times the RADIUS client resends a packet if no ACK is received.

Table 59: show services l2tp radius Output Fields (*continued*)

Field Name	Field Description
<b>Timeout</b>	( <b>servers</b> keyword only) Length of time the client waits for an ACK before retransmission.
<b>Pending Requests</b>	( <b>servers</b> keyword only) Number of client pending authentication or accounting requests.
<b>Maximum Sessions</b>	( <b>servers</b> 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).
<b>Dead Time</b>	( <b>servers</b> keyword only) Interval to wait before retrying a server after it fails to send a response to an authentication or accounting request.
<b>Secret Type</b>	( <b>servers</b> keyword only) Secret type configured on the RADIUS server.
<b>Profile</b>	( <b>servers</b> keyword only) Name of profile configured for the RADIUS server.
<b>Access requests</b>	( <b>statistics</b> keyword only) Number of access requests sent to the server.
<b>Rollover requests</b>	( <b>statistics</b> keyword only) Number of requests coming into the server as a result of the previous server timing out.
<b>Retransmissions</b>	( <b>statistics</b> keyword only) Number of retransmissions.
<b>Access accepts</b>	( <b>statistics</b> keyword only) Number of access accept messages received from the server.
<b>Access rejects</b>	( <b>statistics</b> keyword only) Number of access reject messages received from the server.
<b>Access challenges</b>	( <b>statistics</b> keyword only) Number of access challenges received from the server.
<b>Malformed responses</b>	( <b>statistics</b> 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).
<b>Bad authenticators</b>	( <b>statistics</b> 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.
<b>Requests pending</b>	( <b>statistics</b> keyword only) Number of requests waiting for a response.
<b>Request timeouts</b>	( <b>statistics</b> keyword only) Number of requests that timed out.
<b>Unknown responses</b>	( <b>statistics</b> keyword only) Number of unknown responses. The RADIUS response type in the header is invalid or unsupported.
<b>Packets dropped</b>	( <b>statistics</b> 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 676](#)
- [L2TP Minimum Configuration on page 677](#)
- [clear services l2tp session on page 979](#)

**List of Sample Output**

[show services l2tp session \(LNS on M Series Routers\) on page 1147](#)  
[show services l2tp session \(LNS on MX Series Routers\) on page 1148](#)  
[show services l2tp session \(LAC\) on page 1148](#)  
[show services l2tp session detail \(LAC\) on page 1148](#)  
[show services l2tp session extensive \(LAC\) on page 1148](#)  
[show services l2tp session extensive \(LAC on MX Series Routers\) on page 1148](#)  
[show services l2tp session extensive \(LNS on M Series Routers\) on page 1149](#)  
[show services l2tp session extensive \(LNS on MX Series Routers\) on page 1149](#)  
[show services l2tp session statistics \(MX Series Routers\) on page 1150](#)

**Output Fields** [Table 60 on page 1144](#) lists the output fields for the **show services l2tp session** command. Output fields are listed in the approximate order in which they appear.

**Table 60: 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 60: show services l2tp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Session remote ID</b>	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	All levels
<b>State</b>	State of the L2TP session: <ul style="list-style-type: none"> <li>• <b>Established</b>—Session is operating. This is the only state supported for the LAC.</li> <li>• <b>closed</b>—Session is being closed.</li> <li>• <b>destroyed</b>—Session is being destroyed.</li> <li>• <b>clean-up</b>—Session is being cleaned up.</li> <li>• <b>lns-ic-accept-new</b>—New session is being accepted.</li> <li>• <b>lns-ic-idle</b>—Session has been created and is idle.</li> <li>• <b>lns-ic-reject-new</b>—New session is being rejected.</li> <li>• <b>lns-ic-wait-connect</b>—Session is waiting for the peer's incoming call connected (ICCN) message.</li> </ul>	All levels
<b>Bundle ID</b>	(LNS only) Bundle identifier. Indicates the session is part of a multilink bundle. Sessions that have a blank <b>Bundle</b> 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 <b>show services l2tp multilink extensive</b> command.	All levels
<b>Mode</b>	(LNS) Mode of the interface representing the session: <b>shared</b> or <b>exclusive</b> .  (LAC) Mode of the interface representing the session: <b>shared</b> or <b>dedicated</b> . Only <b>dedicated</b> is currently supported for the LAC.	<b>extensive</b>
<b>Local IP</b>	IP address of local endpoint of the Point-to-Point Protocol (PPP) session.	<b>extensive</b>
<b>Remote IP</b>	IP address of remote endpoint of the PPP session.	<b>extensive</b>
<b>Username</b>	(LNS only) Name of the user logged in to the session.	All levels
<b>Assigned IP address</b>	(LNS only) IP address assigned to remote client.	<b>extensive</b>
<b>Local name</b>	For LNS, name of the LNS instance in which the session was created. For LAC, name of the LAC.	<b>extensive</b>
<b>Remote name</b>	For LNS, name of the LAC from which the session was created. For LAC, name of the LAC instance.	<b>extensive</b>
<b>Local MRU</b>	(LNS only) Maximum receive unit (MRU) setting of the local device, in bytes.	<b>extensive</b>
<b>Remote MRU</b>	(LNS only) MRU setting of the remote device, in bytes.	<b>extensive</b>

Table 60: show services l2tp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Tx speed</b>	<p>Transmit speed of the session conveyed from the LAC to the LNS, in bits per second (bps).</p> <p>Either the initial (<b>initial</b>) line speed or both the initial and current (<b>update</b>) line speeds can be displayed on MX Series routers:</p> <ul style="list-style-type: none"> <li>When connection speed updates are not enabled, then only the initial line speed is displayed.</li> <li>When connection speed updates are enabled, then both the initial and the current speeds are displayed.</li> </ul> <p>When the Tx connect speed method is set to none, the value of zero (0) is displayed.</p>	<b>extensive</b>
<b>Rx speed</b>	<p>Receive speed of the session conveyed from the LAC to the LNS, in bits per second (bps).</p> <p>Either the initial (<b>initial</b>) line speed or both the initial and current (<b>update</b>) line speeds can be displayed on MX Series routers:</p> <ul style="list-style-type: none"> <li>When connection speed updates are not enabled, then only the initial line speed is displayed.</li> <li>When connection speed updates are enabled, then both the initial and the current speeds are displayed.</li> </ul> <p>When the Tx connect speed method is set to none, the value of zero (0) is displayed.</p>	<b>extensive</b>
<b>Bearer type</b>	<p>Type of bearer enabled:</p> <ul style="list-style-type: none"> <li>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).</li> <li>1—Digital access requested.</li> <li>2—Analog access requested.</li> <li>4—Asynchronous Transfer Mode (ATM) bearer support.</li> </ul>	<b>extensive</b>
<b>Framing type</b>	<p>Type of framing enabled:</p> <ul style="list-style-type: none"> <li>1—Synchronous framing</li> <li>2—Asynchronous framing</li> </ul>	<b>extensive</b>
<b>LCP renegotiation</b>	<p>(LNS only) Whether Link Control Protocol (LCP) renegotiation is configured: <b>On</b> or <b>Off</b>.</p>	<b>extensive</b>
<b>Authentication</b>	<p>Type of authentication algorithm used: Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).</p>	<b>extensive</b>
<b>Interface ID</b>	<p>(LNS only) Identifier used to look up the logical interface for this session.</p>	<b>extensive</b>
<b>Interface unit</b>	<p>Logical interface for this session.</p>	All levels
<b>Call serial number</b>	<p>Unique serial number assigned to the call.</p>	<b>extensive</b>

Table 60: show services l2tp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Policer bandwidth</b>	Maximum policer bandwidth configured for this session.	<b>extensive</b>
<b>Policer burst size</b>	Maximum policer burst size configured for this session.	<b>extensive</b>
<b>Firewall filter</b>	Configured firewall filter name.	<b>extensive</b>
<b>Session encapsulation overhead</b>	Overhead allowance configured for this session, in bytes.	<b>extensive</b>
<b>Session cell overhead</b>	Cell overhead activation ( <b>On</b> or <b>Off</b> ).	<b>extensive</b>
<b>Create time</b>	Date and time when the call was created.	<b>extensive</b>
<b>Up time</b>	Length of time elapsed since the call became active, in hours, minutes, and seconds.	<b>extensive</b>
<b>Idle time</b>	Length of time elapsed since the call became idle, in hours, minutes, and seconds.	<b>extensive</b>
<b>Statistics since</b>	Date and time when collection of the following statistics began: <ul style="list-style-type: none"> <li>• <b>Control Tx</b>—Amount of control information transmitted, in packets and bytes.</li> <li>• <b>Control Rx</b>—Amount of control information received, in packets and bytes.</li> <li>• <b>Data Tx</b>—Amount of data transmitted, in packets and bytes.</li> <li>• <b>Data Rx</b>—Amount of data received, in packets and bytes.</li> <li>• <b>Errors Tx</b>—Number of errors transmitted, in packets.</li> <li>• <b>Errors Rx</b>—Number of errors received, in packets.</li> <li>• <b>LCP echo req Tx</b>—Number of LCP echo requests transmitted, in packets.</li> <li>• <b>LCP echo req Rx</b>—Number of LCP echo requests received, in packets.</li> <li>• <b>LCP echo rep Tx</b>—Number of LCP echo responses transmitted, in packets.</li> <li>• <b>LCP echo rep Rx</b>—Number of LCP echo responses received, in packets.</li> <li>• <b>LCP echo Req timeout</b>—Number of LCP echo requests that timed out.</li> <li>• <b>LCP echo Req error</b>—Number of errors received for LCP echo packets.</li> <li>• <b>LCP echo Rep error</b>—Number of errors transmitted for LCP echo packets.</li> </ul>	<b>extensive</b>

## 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@example.com, 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@example.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

<b>Syntax</b>	<code>show services l2tp summary</code> <code>&lt;interface sp-fpc/pic/port&gt;</code> <code>&lt;statistics&gt;</code>
<b>Release Information</b>	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 <b>statistics</b> option introduced in Junos OS Release 13.1.
<b>Description</b>	(M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Display Layer 2 Tunneling Protocol (L2TP) summary information.
<b>Options</b>	<p><b>none</b>—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><b>interface sp-fpc/pic/port</b>—(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><b>statistics</b>—(Optional) Display a summary of control packets and bytes transmitted and received.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">L2TP Services Configuration Overview on page 676</a></li> <li>• <a href="#">L2TP Minimum Configuration on page 677</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show services l2tp summary (LAC on M Series routers) on page 1154</a> <a href="#">show services l2tp summary (LAC on MX Series routers) on page 1154</a> <a href="#">show services l2tp summary (LNS on MX Series routers) on page 1155</a> <a href="#">show services l2tp summary (LNS on M Series routers) on page 1155</a> <a href="#">show services l2tp summary statistics (MX Series routers) on page 1155</a>
<b>Output Fields</b>	Table 61 on page 1151 lists the output fields for the <b>show services l2tp summary</b> command. Output fields are listed in the approximate order in which they appear.

**Table 61: show services l2tp summary Output Fields**

Field Name	Field Description
<b>Administrative state</b>	Administrative state of the tunnel is drain. In this state you cannot configure new sessions, destinations, or tunnels at the LAC or LNS.
<b>Failover within a preference level</b>	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 61: 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 <b>Enabled</b> when a secret is configured in the tunnel profile or on the RADIUS server in the Tunnel-Password attribute [69]. The state is <b>Disabled</b> when the secret is not present. Not displayed for LNS on M Series routers.
Calling number avp	When the state is <b>Enabled</b> , the LAC includes the value of the Calling Number AVP 22 in ICRQ packets sent to the LNS. When the state is <b>Disabled</b> , 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 <b>disable-failover-protocol</b> 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"> <li>• <b>actual</b> This is the default value.</li> <li>• <b>ancp</b></li> <li>• <b>none</b></li> <li>• <b>pppoe-ia-tag</b></li> <li>• <b>static</b></li> </ul>
Rx speed avp when equal	Indicates if the Rx connect speed when equal configuration is <b>enabled</b> or <b>disabled</b> .

Table 61: show services l2tp summary Output Fields (*continued*)

Field Name	Field Description
<b>Tunnel assignment id</b>	<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"> <li>• <b>authentication-id</b>—Name consists of only Tunnel Assignment-Id [82]. This is the default value.</li> <li>• <b>client-server-id</b>—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.</li> </ul>
<b>Tunnel Tx Address Change</b>	<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"> <li>• <b>accept</b>—Accepts change requests for the IP address or UDP port. This is the default action.</li> <li>• <b>ignore</b>—Ignores all change requests.</li> <li>• <b>ignore-ip-address</b>—Ignores change requests for the IP address but accepts them for the UDP port.</li> <li>• <b>ignore-udp-port</b>—Ignores change requests for the UDP port but accepts them for the IP address.</li> </ul>
<b>Min Retransmission Timeout for control packets</b>	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.
<b>Min Retransmission Timeout for control packets</b>	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.
<b>Max Retransmissions for Established Tunnel</b>	Maximum number of times control messages are retransmitted for established tunnels.
<b>Max Retransmissions for Not Established Tunnel</b>	Maximum number of times control messages are retransmitted for tunnels that are not established.
<b>Tunnel Idle Timeout</b>	Period that a tunnel can be inactive—that is, carrying no traffic—before it times out and is torn down.
<b>Destruct Timeout</b>	Period that the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed.
<b>Reassembly Service Set</b>	Indicates active IP reassembly configured for the interface.
<b>Destination Lockout Timeout</b>	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 61: show services l2tp summary Output Fields (*continued*)

Field Name	Field Description
<b>Access Line Information</b>	State of LAC global configuration for forwarding subscriber line information to the LNS, <b>Enabled</b> or <b>Disabled</b> .  Indicates active IP reassembly configured for the interface.
<b>Speed Updates</b>	State of LAC global configuration for including connection speed updates when it forwards subscriber line information to the LNS, <b>Enabled</b> or <b>Disabled</b> .
<b>Destinations</b>	Number of L2TP destinations for the LAC. Not displayed for LNS on M Series routers.
<b>Tunnels</b>	Number of L2TP tunnels established on the router.
<b>Sessions</b>	Number of L2TP sessions established on the router.
<b>Switched sessions</b>	Number of L2TP tunnel-switched sessions established on the router.
<b>Control</b>	Count of L2TP control packets and bytes sent and received.
<b>Data</b>	Count of L2TP data packets and bytes sent and received.
<b>Errors</b>	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

---

<b>Syntax</b>	<pre>show services l2tp tunnel &lt;brief   detail   extensive&gt; &lt;interface sp-fpc/pic/port&gt; &lt;local-gateway gateway-address&gt; &lt;local-gateway-name gateway-name&gt; &lt;local-tunnel-id tunnel-id&gt; &lt;peer-gateway gateway-address&gt; &lt;peer-gateway-name gateway-name&gt; &lt;statistics&gt; &lt;tunnel-group group-name&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	<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; the tunnels may or may not have active sessions.</p>
<b>Options</b>	<p><b>none</b>—Display standard information about all active L2TP tunnels.</p> <p><b>brief   detail   extensive</b>—(Default) Display the specified level of output.</p> <p><b>interface sp-fpc/pic/port</b>—(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><b>local-gateway gateway-address</b>—(Optional) Display L2TP tunnel information for only the specified local gateway address.</p> <p><b>local-gateway-name gateway-name</b>—(Optional) Display L2TP tunnel information for only the specified local gateway name.</p> <p><b>local-tunnel-id tunnel-id</b>—(Optional) Display L2TP tunnel information for only the specified local tunnel identifier.</p> <p><b>peer-gateway gateway-address</b>—(Optional) Display L2TP tunnel information for only the specified peer gateway address.</p> <p><b>peer-gateway-name gateway-name</b>—(Optional) Display L2TP tunnel information for only the specified peer gateway name.</p> <p><b>statistics</b>—(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, <b>brief</b>, <b>detail</b>, or <b>extensive</b>.</p> <p><b>tunnel-group group-name</b>—(Optional) Display L2TP tunnel information for only the specified tunnel group.</p>
<b>Required Privilege Level</b>	view

- Related Documentation**
- [L2TP Services Configuration Overview on page 676](#)
  - [L2TP Minimum Configuration on page 677](#)

- List of Sample Output**
- [show services l2tp tunnel \(LAC\) on page 1159](#)
  - [show services l2tp tunnel detail \(LAC\) on page 1159](#)
  - [show services l2tp tunnel detail \(LAC on MX Series Routers\) on page 1159](#)
  - [show services l2tp tunnel detail \(LNS on MX Series Routers\) on page 1159](#)
  - [show services l2tp tunnel extensive \(LAC\) on page 1160](#)
  - [show services l2tp tunnel extensive \(LNS on M Series Routers\) on page 1160](#)
  - [show services l2tp tunnel extensive \(LNS on MX Series Routers\) on page 1161](#)
  - [show services l2tp tunnel statistics \(MX Series Routers\) on page 1161](#)

- Output Fields** [Table 62 on page 1157](#) lists the output fields for the **show services l2tp tunnel** command. Output fields are listed in the approximate order in which they appear.

**Table 62: show services l2tp tunnel Output Fields**

Field Name	Field Description
<b>Interface</b>	(LNS only) Name of an adaptive services interface.
<b>Tunnel group</b>	(LNS only) Name of a tunnel group.
<b>Local ID</b>	<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>
<b>Remote ID</b>	<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>
<b>Remote IP</b>	IP address of the peer endpoint of the tunnel.
<b>Sessions</b>	Number of L2TP sessions established through the tunnel.

Table 62: show services l2tp tunnel Output Fields (*continued*)

Field Name	Field Description
<b>State</b>	<p>State of the L2TP tunnel:</p> <ul style="list-style-type: none"> <li><b>cc_responder_accept_new</b>—The tunnel has received and accepted the start control connection request (SCCRQ).</li> <li><b>cc_responder_reject_new</b>—The tunnel has received and rejected the SCCRQ.</li> <li><b>cc_responder_idle</b>—The tunnel has just been created.</li> <li><b>cc_responder_wait_ctl_conn</b>—The tunnel has sent the start control connection response (SCCRP) and is waiting for the start control connection connected (SCCCN) message.</li> <li><b>clean-up</b>—The tunnel is being cleaned up.</li> <li><b>closed</b>—The tunnel is being closed.</li> <li><b>destroyed</b>—The tunnel is being destroyed.</li> <li><b>Drain</b>—Creation of new sessions and destinations is disabled for this tunnel.</li> <li><b>Established</b>—The tunnel is operating. This is the only state supported for the LAC.</li> <li><b>Terminate</b>—The tunnel is terminating.</li> <li><b>Unknown</b>—The tunnel is not connected to the router.</li> </ul>
<b>Tunnel Name</b>	(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].
<b>Local IP</b>	IP address of the local endpoint of the tunnel.
<b>Local name</b>	Name used for local tunnel endpoint during tunnel negotiation.
<b>Remote name</b>	Name used for remote tunnel endpoint during tunnel negotiation.
<b>Effective Peer Resync Mechanism</b>	<p>(LAC only) Peer resynchronization mechanism (PRM) in effect for the tunnel:</p> <ul style="list-style-type: none"> <li>Failover protocol</li> <li>Silent failover—Recovery takes place in the failed endpoint only using the proprietary silent failover protocol.</li> </ul>
<b>Nas Port Method</b>	<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"> <li><b>cisco-avp</b>—sends the AVP.</li> <li><b>none</b>—does not send the AVP.</li> </ul>
<b>Tunnel Logical System</b>	Logical system in which the L2TP tunnel is brought up.
<b>Tunnel Routing Instance</b>	Routing instance in which the L2TP tunnel is brought up.
<b>Max sessions</b>	Maximum number of sessions that can be established on this tunnel.
<b>Window size</b>	Number of control messages that can be sent without receipt of an acknowledgment.
<b>Hello interval</b>	Interval between the transmission of hello messages, in seconds.



Table 62: show services l2tp tunnel Output Fields (*continued*)

Field Name	Field Description
<b>Create time</b>	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 <b>Unknown</b> and the <b>Create time</b> value resets.
<b>Up time</b>	Amount of time elapsed since the tunnel became active, in hours, minutes, and seconds.
<b>Idle time</b>	Amount of time elapsed since the tunnel became idle, in hours, minutes, and seconds.
<b>Statistics since</b>	<p>Date and time when collection of the following statistics began:</p> <ul style="list-style-type: none"> <li>• <b>Control Tx</b>—Amount of control information transmitted, in packets and bytes.</li> <li>• <b>Control Rx</b>—Amount of control information received, in packets and bytes.</li> <li>• <b>Data Tx</b>—Amount of data transmitted, in packets and bytes.</li> <li>• <b>Data Rx</b>—Amount of data received, in packets and bytes.</li> <li>• <b>Errors Tx</b>—Number of errors transmitted, in packets.</li> <li>• <b>Errors Rx</b>—Number of errors received, in packets.</li> </ul>

## 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

<b>Syntax</b>	show services l2tp user <brief   detail   extensive   statistics> <user <i>username</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(M10i and M7i routers only) Display a list of active Layer 2 Tunneling Protocol (L2TP) users.
<b>Options</b>	<p><b>none</b>—Display all active L2TP users.</p> <p><b>brief   detail   extensive   statistics</b>—(Optional) Display the specified level of output. Use the <b>statistics</b> option to display L2TP user statistics.</p> <p><b>user <i>username</i></b>—(Optional) Display L2TP user information for only the specified username.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">L2TP Services Configuration Overview on page 676</a></li> <li>• <a href="#">L2TP Minimum Configuration on page 677</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show services l2tp user extensive on page 1164</a>
<b>Output Fields</b>	Table 63 on page 1162 lists the output fields for the <b>show services l2tp user</b> command. Output fields are listed in the approximate order in which they appear.

**Table 63: show services l2tp user Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of an adaptive services interface.
<b>Tunnel group</b>	Name of a tunnel group.
<b>Tunnel local ID</b>	Local identifier of the tunnel, as assigned by the L2TP network server (LNS).
<b>Session local ID</b>	Local identifier of the session, as assigned by the L2TP network server (LNS).
<b>Session remote ID</b>	Remote identifier of the session, as assigned by the L2TP access concentrator (LAC).

Table 63: show services l2tp user Output Fields (*continued*)

Field Name	Field Description
<b>State</b>	State of the L2TP session: <ul style="list-style-type: none"> <li>• <b>Established</b>—The session is operating.</li> <li>• <b>closed</b>—The session is being closed.</li> <li>• <b>destroyed</b>—The session is being destroyed.</li> <li>• <b>clean-up</b>—The session is being cleaned up.</li> <li>• <b>Ins-ic-accept-new</b>—A new session is being accepted.</li> <li>• <b>Ins-ic-idle</b>—The session has been created and is idle.</li> <li>• <b>Ins-ic-reject-new</b>—The new session is being rejected.</li> <li>• <b>Ins-ic-wait-connect</b>—The session is waiting for the peer's incoming call connected (ICCN) message.</li> </ul>
<b>Mode</b>	Mode of the interface representing the session: <b>shared</b> or <b>exclusive</b> .
<b>Local IP</b>	IP address of the local endpoint of the tunnel.
<b>Remote IP</b>	IP address of the peer endpoint of the tunnel.
<b>Username</b>	Name of the user logged in to the session.
<b>Assigned IP address</b>	IP address assigned to remote client.
<b>Local name</b>	Name of the local device.
<b>Remote name</b>	Name of the remote device.
<b>Local MRU</b>	Maximum receive unit (MRU) setting of the local device, in bytes.
<b>Remote MRU</b>	MRU setting of the remote device, in bytes.
<b>Tx speed</b>	Transmit speed of the tunnel session, in bps.
<b>Rx speed</b>	Receive speed of the tunnel session, in bps.
<b>Bearer type</b>	Type of bearer enabled: <ul style="list-style-type: none"> <li>• <b>0</b>—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)</li> <li>• <b>1</b>—Digital access requested</li> <li>• <b>2</b>—Analog access requested</li> <li>• <b>4</b>—Asynchronous Transfer Mode (ATM) bearer support</li> </ul>
<b>Framing type</b>	Type of framing enabled: <ul style="list-style-type: none"> <li>• <b>1</b>—Synchronous framing</li> <li>• <b>2</b>—Asynchronous framing</li> </ul>
<b>LCP renegotiation</b>	Whether Link Control Protocol (LCP) renegotiation is configured: <b>On</b> or <b>Off</b> .

Table 63: show services l2tp user Output Fields (*continued*)

Field Name	Field Description
<b>Authentication</b>	Type of authentication algorithm used: Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).
<b>Interface ID</b>	Name of the logical unit.
<b>Interface unit</b>	Logical unit number.
<b>Call serial number</b>	Unique serial number assigned to the call.
<b>Create time</b>	Date and time when the call was created.
<b>Up time</b>	Amount of time elapsed since the call became active, in hours, minutes, and seconds.
<b>Idle time</b>	Amount of time elapsed since the call became idle, in hours, minutes, and seconds.
<b>Statistics since</b>	<p>Date and time when collection of the following statistics began:</p> <ul style="list-style-type: none"> <li>• <b>Control Tx</b>—Amount of control information transmitted, in packets and bytes.</li> <li>• <b>Control Rx</b>—Amount of control information received, in packets and bytes.</li> <li>• <b>Data Tx</b>—Amount of data transmitted, in packets and bytes.</li> <li>• <b>Data Rx</b>—Amount of data received, in packets and bytes.</li> <li>• <b>Errors Tx</b>—Number of errors transmitted, in packets.</li> <li>• <b>Errors Rx</b>—Number of errors received, in packets.</li> </ul>

## 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@example.com, 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 deterministic-nat internal-host

<b>Syntax</b>	<code>show services nat deterministic-nat internal-host</code> <i>nat-address</i> <i>nat-port</i>
<b>Release Information</b>	Command introduced in Junos OS Release 12.1.
<b>Description</b>	This command prints the internal host address and algorithmically determined port ranges for the specified NAT IP address and port number. The results are calculated on the PIC and the results are sent to RE.
<b>Options</b>	<i>nat-address</i> —NAT address of the internal host.  <i>nat-port</i> —NAT port of the internal host.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services nat deterministic-nat internal-host on page 1166</a>
<b>Output Fields</b>	<a href="#">Table 64 on page 1166</a> lists the output fields for the <code>show services nat mapping</code> command. Output fields are listed in the approximate order in which they appear.

**Table 64: show services nat deterministic-nat internal-host Output Fields**

Field Name	Field Description
Interface	Name of a service interface.
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.
Internal Host	Private IP address of a subscriber on the access network.
NAT IP address	NAT public IP address
NAT Port Start	Lowest port number in range of assigned ports.
NAT Port End	Highest port number in range of assigned ports.

## Sample Output

### show services nat deterministic-nat internal-host

```

user@host> show services nat deterministic-nat internal-host 203.0.113.1 2000
Service set: ss1
Interface: sp-2/0/0
NAT pool: pool1
Internal Host: 192.0.2.4, NAT IP Address: 203.0.113.1, NAT Port Start: 1792, NAT
Port End: 2047

```





## show services nat deterministic-nat nat-port-block

<b>Syntax</b>	<code>show services nat deterministic-nat nat-port-block</code> <i>nat-address</i> <i>nat-port</i>
<b>Release Information</b>	Command introduced in Junos OS Release 12.1.
<b>Description</b>	Display the translated NAT address and port ranges for the given internal host.
<b>Options</b>	<i>internal-host</i> —IP address of the internal host.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">run show services nat deterministic-nat nat-port-block on page 1168</a>
<b>Output Fields</b>	<a href="#">Table 65 on page 1168</a> lists the output fields for the <b>show services nat deterministic-nat nat-port-block</b> command. Output fields are listed in the approximate order in which they appear.

**Table 65: show services nat deterministic-nat nat-port-block Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of a service interface.
<b>Service set</b>	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.
<b>Internal Host</b>	Private IP address of a subscriber on the access network.
<b>NAT IP address</b>	NAT public IP address
<b>NAT Port Start</b>	Lowest port number in range of assigned ports.
<b>NAT Port End</b>	Highest port number in range of assigned ports.

## Sample Output

### run show services nat deterministic-nat nat-port-block

```

user@host> show services nat deterministic-nat nat-port-block 128.1.1.1
root@host# run show services nat deterministic-nat nat-port-block 128.1.1.1
Service set: ss1
Interface: sp-2/0/0
NAT pool: pool1
Internal Host: 128.1.1.1, NAT IP Address: 32.32.32.1, NAT Port Start: 1024, NAT
Port End: 1279

```

## show services nat ipv6-multicast-interfaces

<b>Syntax</b>	show services nat ipv6-multicast-interfaces
<b>Release Information</b>	Command introduced in Junos OS Release 8.5.
<b>Description</b>	Displays a list of interfaces enabled for IPv6 mutlicast.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services nat ipv6-multicast-interfaces on page 1169</a>
<b>Output Fields</b>	<a href="#">Table 66 on page 1169</a> lists the output fields for the <b>show services nat ipv6-multicast-interfaces</b> command. Output fields are listed in the approximate order in which they appear.

Table 66: show services nat ipv6-multicast-interfaces Output Fields

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of a service interface.	All levels
<b>Admin State</b>	Configured IPv6 multicast capability of an interface ,	All levels
<b>Operational State</b>	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

<b>Syntax</b>	<pre>show services nat mappings &lt;brief   detail   summary&gt; &lt;nptv6 (ipv6-address   external ipv6-address   internal ipv6-address)&gt; &lt;pool-name&gt; &lt;address-pooling-paired   endpoint-independent   pcp&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 10.1.</p> <p><b>summary</b> option introduced in Junos OS Release 11.1.</p> <p><b>address-pooling paired</b> option introduced in Junos OS Release 13.2.</p> <p><b>endpoint-independent</b> option introduced in Junos OS Release 13.2.</p> <p><b>pcp</b> option introduced in Junos OS Release 13.2.</p> <p><b>nptv6</b> option introduced in Junos OS Release 15.1.</p>
<b>Description</b>	Display information about Network Address Translation (NAT) address, port, and port control protocol (PCP) mappings.
<b>Options</b>	<p><b>none</b>—Display standard information about all NAT pools.</p> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output.</p> <p><b>nptv6</b>—(Optional) Display information about the network prefix translation for IPv6 traffic.</p> <p><b>ipv6-address</b>—(Optional) Display the network prefix translation details for the specified IPv6 address.</p> <p><b>external</b>—(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><b>internal</b>—(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><b>pool-name</b>—(Optional) Display detailed information about a specific NAT pool. Used only with detail level output.</p> <p><b>address-pooling-paired</b>—(Optional) Display only information about address-pooling paired mappings.</p> <p><b>endpoint-independent</b>—(Optional) Display only information about endpoint-independent mappings.</p> <p><b>pcp</b>—(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 1173](#)
- [show services nat mapping detail on page 1174](#)
- [show services nat mappings pool-name on page 1174](#)
- [show services nat mappings summary on page 1174](#)
- [show services nat mappings address-pooling-paired on page 1174](#)
- [show services nat mappings address-pooling-paired \(mapping of active B4 for a subscriber\) on page 1174](#)
- [show services nat mappings endpoint-independent on page 1175](#)
- [show services nat mappings pcp on page 1175](#)
- [show services nat mappings nptv6 internal on page 1175](#)
- [show services nat mappings nptv6 external on page 1175](#)

**Output Fields** [Table 67 on page 1172](#) lists the output fields for the **show services nat mappings** command. Output fields are listed in the approximate order in which they appear.

**Table 67: 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 67: 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"> <li>• <b>ACTIVE</b>—Indicates that the entry is active and in use.</li> <li>• <b>TIMEOUT</b>—Indicates that the mapping is not in use. After the <b>mapping-timeout</b>, configured at the <b>[edit services nat pool pool-name]</b> hierarchy level, lapses, the mapping is deleted. This field also displays the number of seconds after which the timeout occurs.</li> </ul>	
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

---

<b>Syntax</b>	<code>show services nat pool</code> <code>&lt;brief   detail&gt;</code> <code>&lt;pool-name&gt;</code> <code>pgcp &lt;ports-per-session   remotely-controlled&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. <code>pgcp</code> option added in Junos OS Release 8.5.
<b>Description</b>	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.

---

<b>Options</b>	<code>none</code> —Display standard information about all NAT pools.
	<code>brief   detail</code> —(Optional) Display the specified level of output.
	<code>pool-name</code> —(Optional) Display information about the specified NAT pool.
	<code>pgcp</code> —(Optional) Display information about a NAT pool that is exclusive to the BGF.
	<code>ports-per-session</code> —(Optional) Display the number of ports allocated per session from the NAT pool.
	<code>remotely-controlled</code> —(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 1178](#)  
[show services nat pool detail on page 1179](#)  
[show services nat pool for Secured Port Block Allocation on page 1179](#)  
[show services nat pool detail for Deterministic Port Block Allocation on page 1179](#)  
[show services nat pool for Deterministic Port Block Allocation on page 1180](#)  
[show services nat pool detail for Port Block Allocation on page 1180](#)

**Output Fields** [Table 68 on page 1177](#) lists the output fields for the **show services nat pool** command. Output fields are listed in the approximate order in which they appear.

**Table 68: show services nat pool Output Fields**

Field Name	Field Description	Level of Output
<b>DetNat subscriber exceeded port limits</b>	The number of times a subscriber exceeded its port limits for a NAT pool that uses deterministic port block allocation.	All levels.
<b>MAX number of port blocks used</b>	The maximum number of port blocks used.	All levels.
<b>Port block memory allocation errors</b>	The number of memory allocation errors for port blocks.	All levels.
<b>Current number of port blocks in use</b>	Current count of the port blocks that are being used.	
<b>Unique pool users</b>	The number of different users of the NAT pools.	All levels.
<b>Port block allocation errors</b>	The consolidated number of port block allocation errors.	All levels.
<b>Port blocks limit exceeded errors</b>	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.
<b>Interface</b>	Name of an adaptive services interface.	All levels
<b>Service set</b>	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
<b>NAT pool</b>	Name of the Network Address Translation pool.	All levels
<b>Type or Translation type</b>	Address translation type: <b>basic-nat-pt</b> , <b>basic-nat44</b> , <b>basic-nat66</b> , <b>deterministic-napt44</b> , <b>dnat-44</b> , <b>dynamic-nat44</b> , <b>napt44</b> , <b>napt-66</b> , <b>napt-pt</b> , <b>stateful-nat64</b> , <b>twice-basic-nat-44</b> , <b>twice-dynamic-nat-44</b> , <b>twice-dynamic-napt-44</b> .	All levels
<b>Address or Address range</b>	IPv4 address range of the pool.	All levels

Table 68: show services nat pool Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Configured port range</b>	The range of ports configured to be used for NAT pool.	<b>detail</b>
<b>Preserve range enabled</b>	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"> <li>• <b>Is active</b>—Preservation of port range is enabled.</li> <li>• <b>Not active</b>—Preservation of port range is not enabled.</li> </ul>	<b>detail</b>
<b>Port or Port range</b>	Port range of the pool. Applicable only for dynamic NAT pools. Not displayed for static NAT pools.	All levels
<b>Ports used' or Ports in use</b>	Number of ports allocated in this pool with this name. Applicable only for dynamic NAT pools. Not displayed for static NAT pools.	All levels
<b>Port block type</b>	Type of port block allocation: secured or deterministic	All levels
<b>Out of port errors</b>	Number of port allocation errors. Applicable only for dynamic NAT pools. Not displayed for static NAT pools.	<b>detail</b>
<b>Max ports used</b>	Maximum number of ports used. Applicable only for dynamic NAT pools. Not displayed for static NAT pools.	<b>detail</b>
<b>Addresses in use</b>	Number of addresses in use for dynamic source address NAT pools.	<b>detail</b>
<b>Out of Port Errors</b>	No more ports available to allocate.	Detail
<b>Max Ports Used</b>	The maximum number of ports in use at any time since the services PIC was started.	Detail
<b>AP-P out of port errors</b>	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
<b>Current EIF Inbound flows count</b>	Current count of EIF inbound flows, including all EIF flows per pool.	
<b>EIF flow limit exceeded drops</b>	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               1024-63487  0
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

<b>Syntax</b>	show services pcsp statistics
<b>Release Information</b>	Command introduced in Junos OS Release 13.2
<b>Description</b>	Display information PCP mappings.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services pcsp statistics pcsp on page 1182</a>
<b>Output Fields</b>	<a href="#">Table 69 on page 1181</a> lists the output fields for the <b>show services pcsp statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 69: 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 <b>prefer-fail</b> and <b>third-party</b> .
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 69: 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 <b>third-party</b> 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

<b>Syntax</b>	show services service-sets cpu-usage <interface <i>interface-name</i> > <service-set <i>service-set-name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display service set CPU usage as a percentage. The command is supported only on Adaptive Services PICs (SP PICs).
<b>Options</b>	<p><b>none</b>—Display CPU usage for all adaptive services interfaces and service sets.</p> <p><b>interface <i>interface-name</i></b>—(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><b>service-set <i>service-set-name</i></b>—(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>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services service-sets cpu-usage on page 1184</a>
<b>Output Fields</b>	<a href="#">Table 70 on page 1184</a> lists the output fields for the <b>show services service-sets cpu-usage</b> command. Output fields are listed in the approximate order in which they appear.

**Table 70: show services service-sets cpu-usage Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of an adaptive services interface
<b>Service set (system category)</b>	Name of the CPU usage category: <ul style="list-style-type: none"> <li>• idp_recommended—Name of the service sets (displays all the service sets attached to the service PICs)</li> <li>• Idle</li> <li>• System</li> <li>• Receive</li> <li>• Transmit</li> </ul>
<b>CPU utilization %</b>	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 1187](#)  
[show services service-sets memory-usage zone on page 1187](#)  
[show services service-sets memory-usage interface on page 1187](#)

**Output Fields** [Table 71 on page 1186](#) 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 71: show services service-sets memory-usage Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of an adaptive services interface
<b>Service set</b>	Name of a service set
<b>Bytes Used</b>	Number of bytes of memory being used

Table 71: show services service-sets memory-usage Output Fields (*continued*)

Field Name	Field Description
<b>Memory zone</b>	<p>Memory zone in which the adaptive services interface is currently operating:</p> <ul style="list-style-type: none"> <li>• <b>Green</b>—All new flows are allowed.</li> <li>• <b>Yellow</b>—Unused memory is reclaimed. All new flows are allowed.</li> <li>• <b>Orange</b>—New flows are allowed only for service sets that are using less than their equal share of memory.</li> <li>• <b>Red</b>—No new flows are allowed.</li> </ul>

## 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 integrity-drops

<b>Syntax</b>	show services service-sets statistics integrity-drops <interface <i>interface-name</i> > <service-set <i>service-set-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 13.1
<b>Description</b>	Display integrity-drops statistics for one adaptive services interface, for all adaptive services interfaces, or for one service-set. You can configure use the output of this command to verify the packet header for anomalies in IP, TCP, UDP, and IGMP information and to examine any anomalies and errors.
<b>Options</b>	<p><b>none</b>—Display integrity-drops statistics for all configured adaptive service interfaces/ service-set.</p> <p><b>service-set <i>service-set-name</i></b>—(Optional) Display integrity-drops statistics for the specified service-set</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display integrity-drops statistics for the specified adaptive services interface.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear services service-sets statistics integrity-drops on page 1000</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show services service-sets statistics integrity-drops on page 1190</a>
<b>Output Fields</b>	Table 72 on page 1188 lists the output fields for the <b>show services service-sets statistics integrity-drops</b> command. Output fields are listed in the approximate order in which they appear.

**Table 72: show services service-sets statistics integrity-drops Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of an adaptive services interface.
<b>Service set</b>	Name of a service set.
<b>Errors</b>	<p>Total errors, categorized by protocol:</p> <ul style="list-style-type: none"> <li><b>IP</b>—Total IP version 4 errors.</li> <li><b>TCP</b>—Total Transmission Control Protocol (TCP) errors.</li> <li><b>UDP</b>—Total User Datagram Protocol (UDP) errors.</li> <li><b>ICMP</b>—Total Internet Control Message Protocol (ICMP) errors.</li> </ul>

**Table 72: show services service-sets integrity-drops Output Fields (continued)**

Field Name	Field Description
IP Errors	<p>IPv4 errors:</p> <ul style="list-style-type: none"> <li>• <b>IP packet length inconsistencies</b>—IP packet length does not match the Layer 2 reported length.</li> <li>• <b>Minimum IP header length check failures</b>—Minimum IP header length is 20 bytes. The received packet contains less than 20 bytes.</li> <li>• <b>Reassembled packet exceeds maximum IP length</b>—After fragment reassembly, the reassembled IP packet length exceeds 65,535.</li> <li>• <b>Illegal source address 0</b>—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.</li> <li>• <b>Illegal destination address</b> —Destination address is not a valid address. The address is reserved.</li> <li>• <b>TTL zero errors</b>—Received packet had a time-to-live (TTL) value of 0.</li> <li>• <b>Illegal IP protocol number 0 or 255</b>—IP protocol is 0 or 255.</li> <li>• <b>Land attack</b>—IP source address is the same as the destination address.</li> <li>• <b>Non-IP packets</b>—Packet did not conform to the IP standard.</li> <li>• <b>IP option</b>—Packet dropped because of a nonallowed IP option.</li> <li>• <b>Non-IPv4 packets</b>—Packet was not of the IPv4 type.</li> <li>• <b>Non-IPv6 packets</b>—Packet was not of the IPv6 type.</li> <li>• <b>Bad checksum</b>—Packet had an invalid IP checksum.</li> <li>• <b>Illegal IP fragment length</b>—Illegal fragment length. All fragments (other than the last fragment) must have a length that is a multiple of 8 bytes.</li> <li>• <b>IP fragment overlap</b>—Fragments have overlapping fragment offsets.</li> <li>• <b>IP fragment limit exceeded:</b> —Fragments dropped because the configured number of allowed fragments for a packet was exceeded.</li> <li>• <b>IP fragment reassembly timeout</b>—Some of the fragments for an IP packet were not received in time, and the reassembly handler dropped partial fragments. Whenever a fragment is received, it is maintained in a chain until all other fragments are received. If other fragments do not arrive within the configured value of reassembly-timeout, this packet is dropped and the value of the counter shown in this field is incremented. If other fragments arrive in time but the total number of fragments is more than the configured value of fragment-limit, all the fragments (of this packet) are dropped and the value of the counter shown in this field is incremented.</li> <li>• <b>Unknown:</b> —Unknown fragments.</li> </ul>

Table 72: show services service-sets integrity-drops Output Fields (*continued*)

Field Name	Field Description
TCP Errors	<p>TCP protocol errors:</p> <ul style="list-style-type: none"> <li>• <b>TCP header length inconsistencies</b>—Minimum TCP header length is 20 bytes, and the IP packet received does not contain at least 20 bytes.</li> <li>• <b>Source or destination port number is zero</b>—TCP source or destination port is zero.</li> <li>• <b>Illegal sequence number, flags combination</b>—Dropped because of TCP errors, such as an illegal sequence number, which causes an illogical combination of flags to be set.</li> </ul>
UDP Errors	<p>UDP protocol errors:</p> <ul style="list-style-type: none"> <li>• <b>IP data length less than minimum UDP header length (8 bytes)</b>—Minimum UDP header length is 8 bytes. The received IP packets contain less than 8 bytes.</li> <li>• <b>Source or destination port is zero</b>—UDP source or destination port is 0.</li> </ul>
ICMP Errors	<p>ICMP protocol errors:</p> <ul style="list-style-type: none"> <li>• <b>IP data length less than minimum ICMP header length (8 bytes)</b>—ICMP header length is 8 bytes. This counter is incremented when received IP packets contain less than 8 bytes.</li> <li>• <b>ICMP error length inconsistencies</b>—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.</li> </ul>

## Sample Output

### show services service-sets statistics integrity-drops

```

user@host> show services service-sets statistics integrity-drops
Interface: ms-1/0/0
Service set: sset1
Errors:
  IP: 0, TCP: 0
  UDP: 0, ICMP: 0
IP errors:
  IP packet length inconsistencies: 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
  Non-IPv6 packets: 0
  Bad checksum: 0
  Illegal IP fragment length: 0
  IP fragment overlap: 0
  IP fragment limit exceeded: 0
  IP fragment reassembly timeout: 0
  Unknown: 0

```



```
TCP errors:
  TCP header length inconsistencies: 0
  Source or destination port number is zero: 0
  Illegal sequence number and flags combinations: 0
UDP errors:
  IP data length less than minimum UDP header length (8 bytes): 0
  Source or destination port number is zero: 0
ICMP errors:
  IP data length less than minimum ICMP header length (8 bytes): 0
  ICMP error length inconsistencies: 0
```

## show services service-sets statistics packet-drops

<b>Syntax</b>	show services service-sets statistics packet-drops <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 7.4.
<b>Description</b>	Display the number of dropped packets for service sets exceeding CPU limits or memory limits.
<b>Options</b>	<p><b>none</b>—Display the number of dropped service sets packets for all adaptive services interfaces.</p> <p><b>interface <i>interface-name</i></b>—(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>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear services flow-collector statistics</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show services service-sets statistics packet-drops interface on page 1192</a>
<b>Output Fields</b>	<a href="#">Table 72 on page 1188</a> lists the output fields for the <b>show services service-sets packet-drops</b> command. Output fields are listed in the approximate order in which they appear.

**Table 73: 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

<b>Syntax</b>	show services service-sets statistics syslog <interface <i>interface-name</i> > <service-set <i>service-set-name</i> > <brief   detail>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1.
<b>Description</b>	Display the system log statistics with optional filtering by interface and service set name..
<b>Options</b>	<p><b>none</b>—Display the system log statistics for all services interfaces and all service sets.</p> <p><b>brief</b>—(Default) Display abbreviated system log statistics.</p> <p><b>detail</b>—Display detailed system log statistics.</p> <p><b>interface <i>interface-name</i></b>—(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><b>service-set <i>service-set name</i></b>—(Optional) Display the system log statistics for a specific named service-set.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear services service-sets statistics syslog on page 1002</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show services service-sets statistics syslog brief on page 1197</a> <a href="#">show services service-sets statistics syslog detail on page 1198</a>
<b>Output Fields</b>	Table 74 on page 1194 lists the output fields for the <b>show services service-sets statistics syslog</b> command. Output fields are listed in the approximate order in which they appear.

Table 74: show services service-sets statistics syslog Output Fields

Field Name	Field Description	Level
Interface	Name of a services interface.	all
Rate limit	Maximum number of messages per second written to the interface's system log.	all
Sent	Number of messages sent that are not associated with a service set.	all
Dropped	Number of messages dropped that are not associated with a service set.	all
Service-set		

Table 74: show services service-sets statistics syslog Output Fields (*continued*)

Field Name	Field Description	Level
<b>Service set</b>	Name of a service set.	all
<b>Sent</b>	Number of messages sent.	all
<b>Dropped</b>	Number of messages dropped.	all
<b>Session open logs</b>	<p>The following information is displayed for system log messages for session open events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> <li>• <b>Sent</b>—Number of messages sent.</li> <li>• <b>Dropped</b>—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> <li>• <b>low priority</b>—Priority of the message was too low for the message to be sent.</li> <li>• <b>no class set</b>—Specific classes of event messages were configured and this class was not selected.</li> <li>• <b>above rate limit</b>—Maximum number of system log messages per second was exceeded.</li> </ul> </li> </ul>	<b>detail</b>
<b>Session close logs</b>	<p>The following information is displayed for system log messages for session close events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> <li>• <b>Sent</b>—Number of messages sent.</li> <li>• <b>Dropped</b>—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> <li>• <b>low priority</b>—Priority of the message was too low for the message to be sent.</li> <li>• <b>no class set</b>—Specific classes of event messages were configured and this class was not selected.</li> <li>• <b>above rate limit</b>—Maximum number of system log messages per second was exceeded.</li> </ul> </li> </ul>	<b>detail</b>
<b>Packet logs</b>	<p>The following information is displayed for system log messages for packet events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> <li>• <b>Sent</b>—Number of messages sent.</li> <li>• <b>Dropped</b>—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> <li>• <b>low priority</b>—Priority of the message was too low for the message to be sent.</li> <li>• <b>no class set</b>—Specific classes of event messages were configured and this class was not selected.</li> <li>• <b>above rate limit</b>—Maximum number of system log messages per second was exceeded.</li> </ul> </li> </ul>	<b>detail</b>

Table 74: show services service-sets statistics syslog Output Fields (*continued*)

Field Name	Field Description	Level
<b>Stateful firewall logs</b>	<p>The following information is displayed for system log messages for stateful firewall events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> <li>• <b>Sent</b>—Number of messages sent.</li> <li>• <b>Dropped</b>—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> <li>• <b>low priority</b>—Priority of the message was too low for the message to be sent.</li> <li>• <b>no class set</b>—Specific classes of event messages were configured and this class was not selected.</li> <li>• <b>above rate limit</b>—Maximum number of system log messages per second was exceeded.</li> </ul> </li> </ul>	<b>detail</b>
<b>ALG logs</b>	<p>The following information is displayed for system log messages for ALG events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> <li>• <b>Sent</b>—Number of messages sent.</li> <li>• <b>Dropped</b>—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> <li>• <b>low priority</b>—Priority of the message was too low for the message to be sent.</li> <li>• <b>no class set</b>—Specific classes of event messages were configured and this class was not selected.</li> <li>• <b>above rate limit</b>—Maximum number of system log messages per second was exceeded.</li> </ul> </li> </ul>	<b>detail</b>
<b>NAT logs</b>	<p>The following information is displayed for system log messages for NAT events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> <li>• <b>Sent</b>—Number of messages sent.</li> <li>• <b>Dropped</b>—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> <li>• <b>low priority</b>—Priority of the message was too low for the message to be sent.</li> <li>• <b>no class set</b>—Specific classes of event messages were configured and this class was not selected.</li> <li>• <b>above rate limit</b>—Maximum number of system log messages per second was exceeded.</li> </ul> </li> </ul>	<b>detail</b>

Table 74: show services service-sets statistics syslog Output Fields (*continued*)

Field Name	Field Description	Level
IDS logs	<p>The following information is displayed for system log messages for IDS events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> <li>• <b>Sent</b>—Number of messages sent.</li> <li>• <b>Dropped</b>—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> <li>• <b>low priority</b>—Priority of the message was too low for the message to be sent.</li> <li>• <b>no class set</b>—Specific classes of event messages were configured and this class was not selected.</li> <li>• <b>above rate limit</b>—Maximum number of system log messages per second was exceeded.</li> </ul> </li> </ul>	<b>detail</b>
Other logs	<p>The following information is displayed for system log messages for other types of events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> <li>• <b>Sent</b>—Number of messages sent.</li> <li>• <b>Dropped</b>—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> <li>• <b>low priority</b>—Priority of the message was too low for the message to be sent.</li> <li>• <b>no class set</b>—Specific classes of event messages were configured and this class was not selected.</li> <li>• <b>above rate limit</b>—Maximum number of system log messages per second was exceeded.</li> </ul> </li> </ul>	<b>detail</b>

## Sample Output

### show services service-sets statistics syslog brief

```

user@host> show services service-sets statistics syslog brief
Interface: sp-1/1/0
  Rate limit: 200000
  Sent: 0
  Dropped: 0
  Service-set: sset-sfw-sp1
    Sent: 20
    Dropped: 3488
  Service-set: sset-nat-sp1
    Sent: 18
    Dropped: 91
Interface: sp-1/2/0
  Rate limit: 15000
  Sent: 0
  Dropped: 0
  Service-set: sset-sfw-sp2
    Sent: 210
    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
  Rate limit: 10
  Sent: 0
  Dropped: 0
  Service-set: sset-sfw
    Sent: 0
    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

<b>Syntax</b>	show services service-sets statistics tcp-mss <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.5.
<b>Description</b>	(M Series and T Series routers only) Display TCP maximum segment size (MSS) statistics for service sets.
<b>Options</b>	<p><b>none</b>—Display service set TCP MSS information for all adaptive services interfaces.</p> <p><b>interface <i>interface-name</i></b>—(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>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services service-sets statistics tcp-mss on page 1199</a>
<b>Output Fields</b>	Table 75 on page 1199 lists the output fields for the <b>show services service-sets statistics tcp-mss</b> command. Output fields are listed in the approximate order in which they appear.

**Table 75: show services service-sets statistics tcp-mss Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of the adaptive services interface.
<b>Service Set</b>	Name of the configured service set.
<b>SYN Received</b>	Number of TCP SYN packets received.
<b>SYN Modified</b>	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

<b>Syntax</b>	show services service-sets summary <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display service set summary information.
<b>Options</b>	<p><b>none</b>—Display service set summary information for all adaptive services interfaces.</p> <p><b>interface <i>interface-name</i></b>—(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>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services service-sets summary on page 1200</a> <a href="#">show services service-sets summary interface on page 1201</a>
<b>Output Fields</b>	Table 76 on page 1200 lists the output fields for the <b>show services service-sets summary</b> command. Output fields are listed in the approximate order in which they appear.

**Table 76: show services service-sets summary Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of an adaptive services interface
<b>Service type</b>	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)
<b>Service sets configured</b>	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
<b>Bytes used</b>	Bytes used by a particular service or all services
<b>Policy bytes used</b>	Policy bytes used by a particular service or all services
<b>CPU utilization</b>	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
CRTP	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 sessions

---

**Syntax**    show services sessions  
              <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*>  
              <utilization>

**Release Information**    Command introduced in Junos OS Release 10.4.

**Description**    Display session information.



**NOTE:** On MX Series routers (with interchassis redundancy configured), the idle timeout for every flow is displayed in the **show services session extensive** and **show services flows extensive** commands.

---

**Options**    **none**—Display standard information about all sessions.

**brief | extensive | terse**—(Optional) Display the specified level of output.

**application-protocol**—(Optional) Display information about one of the following application protocols:

- **bootp**—Bootstrap protocols
- **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**—Remote Execution Protocol
- **ftp**—File Transfer Protocol
- **h323**—H.323
- **icmp**—ICMP
- **icmpv6**—ICMPv6
- **iiop**—Internet Inter-ORB Protocol
- **ip**—IP

- **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
- **rsh**—Remote Shell
- **sip**—Session Initiation Protocol
- **shell**—Shell
- **snmp**—SNMP
- **sql**—SQLNet
- **talk**—Talk Program
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame



**NOTE:** You can use the **none** option with the **show services sessions count application-protocol** command to display information about sessions other than ALG sessions.

**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**. On J Series routers, *interface-name* is **ms-pim/O/port**.

**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.

**utilization**—(Optional) Display statistical details about session utilization.

**Required Privilege  
Level**

view

**List of Sample Output**

[show services sessions on page 1206](#)  
[show services sessions brief on page 1206](#)  
[show services sessions extensive on page 1206](#)  
[show services sessions terse on page 1206](#)  
[show services sessions application-protocol on page 1206](#)  
[show services sessions count on page 1208](#)  
[show services sessions destination-port on page 1208](#)  
[show services sessions destination-prefix on page 1208](#)  
[show services sessions interface on page 1208](#)  
[show services sessions protocol on page 1209](#)  
[show services sessions service-set on page 1209](#)  
[show services sessions source-port on page 1209](#)  
[show services sessions source-prefix on page 1209](#)

**Output Fields**

[Table 77 on page 1205](#) lists the output fields for the **show services sessions** command. Output fields are listed in the approximate order in which they appear.

Table 77: show services sessions Output Fields

Field Name	Field Description
<b>Interface</b>	Name of the interface.
<b>Session ID</b>	Session ID that uniquely identifies the session.
<b>ALG</b>	Name of the application.
<b>Flags</b>	Session flag for the ALG: <ul style="list-style-type: none"> <li>• <b>0x1</b>—Found an existing session.</li> <li>• <b>0x2</b>—Reached session or flow limit.</li> <li>• <b>0x3</b>—No memory available for new sessions.</li> <li>• <b>0x4</b>—No free session ID available.</li> <li>• <b>0x0000</b>—No session ID found.</li> </ul>
<b>IP Action</b>	Flag indicating whether IP action has been set for the session..
<b>Offload</b>	Flag indicating whether the session has been offloaded to the Packet Forwarding Engine.
<b>Asymmetric</b>	Flag indicating whether the session is uni-directional.
<b>Service set</b>	Name of a service set. Individual empty service sets are not displayed.
<b>Sessions Count</b>	Number of sessions.
<b>Flow or Flow Prot</b>	Protocol used for this session.
<b>Source</b>	Source prefix of the flow in the format <i>source-prefix:port</i> . For ICMP flows, port information is not displayed.
<b>Dest</b>	Destination prefix of the flow. For ICMP flows, port information is not displayed.
<b>State</b>	Status of the flow: <ul style="list-style-type: none"> <li>• <b>Drop</b>—Drop all packets in the flow without response.</li> <li>• <b>Forward</b>—Forward the packet in the flow without looking at it.</li> <li>• <b>Reject</b>—Drop all packets in the flow with response.</li> <li>• <b>Watch</b>—Inspect packets in the flow.</li> <li>• <b>Bypass</b>—Bypass packets in the flow.</li> <li>• <b>Unknown</b>—Unknown flow status.</li> </ul>
<b>Packet Direction</b>	Direction of the flow: ingress ( <b>I</b> ), egress ( <b>O</b> ) or unknown.
<b>Frm count</b>	Number of frames in the flow.

## Sample Output

### show services sessions

```

user@host> show services sessions
ms-2/0/0
Session: 293, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    10.10.10.2:43677 -> 10.20.20.1:53 Forward I      1
UDP    10.20.20.1:53   -> 1.1.1.1:43677 Forward 0      1
Session: 53, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    10.10.10.2:37494 -> 10.20.20.1:53 Forward I      1
UDP    10.20.20.1:53   -> 10.11.11.11:37494 Forward 0      1
Session: 66, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    10.10.10.2:48161 -> 10.20.20.1:53 Forward I      1
UDP    10.20.20.1:53   -> 10.11.11.11:48161 Forward 0      1
Session: 17, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    10.10.10.2:38908 -> 10.20.20.1:53 Forward I      1
UDP    10.20.20.1:53   -> 10.11.11.11:38908 Forward 0      1
Session: 42, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    10.10.10.2:58189 -> 10.20.20.1:53 Forward I      1
UDP    10.20.20.1:53   -> 10.11.11.11:58189 Forward 0      1

```

### show services sessions brief

The output for the **show services flows brief** command is identical to that for the **show services sessions** command. For sample output, see [show services sessions on page 1206](#).

### show services sessions extensive

```

user@host> show services sessions extensive
ms-0/1/0
Session: 2, ALG: 0, Flags: 0x0080, IP Action: no, Offload: no
NAT Plugin Data:
  NAT Action: Translation Type - DYNAMIC NAT44
  NAT source 3.1.1.2 -> 10.10.10.127
TCP    3.1.1.2:52145 -> 4.1.1.2:23 Forward I      22
  Byte count: 1483
  Flow role: Unknown, Timeout: 0
TCP    4.1.1.2:23 -> 10.10.10.127:52145 Forward 0      18
  Byte count: 2712
  Flow role: Unknown, Timeout: 0

```

### show services sessions terse

```

user@host> show services sessions terse
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP    10.2.2.2:52138 -> 10.1.1.2:21 Forward I      33
TCP    10.1.1.2:21 -> 10.2.2.2:52138 Forward 0      31

```

### show services sessions application-protocol

This command has the same output for the rpc, dce-rpc, rpc-portmap and dce-rpc-portmap ALGs.

```

user@host> show services sessions application-protocol dce-rpc
Interface name: ms-1/1/0
Session: 8, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:1019 -> 192.168.203.194:2049 Forward I      4

```



```

UDP    192.168.203.194:2049 ->192.168.203.198:1019 Forward 0      4
Session: 7, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:954 ->192.168.203.194:613 Forward I      1
UDP    192.168.203.194:613 ->192.168.203.198:954 Forward 0      1
Session: 6, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:53836 ->192.168.203.194:613 Forward I     1
UDP    192.168.203.194:613 ->192.168.203.198:53836 Forward 0    1
Session: 5, ALG: portmapper, Flags: 0x1000, IP Action: no, Offload: no
UDP    192.168.203.198:59813 ->192.168.203.194:111 Forward I    1
UDP    192.168.203.194:111 ->192.168.203.198:59813 Forward 0    1
Session: 4, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:36595 ->192.168.203.194:2049 Forward I    1
UDP    192.168.203.194:2049 ->192.168.203.198:36595 Forward 0    1
Session: 3, ALG: portmapper, Flags: 0x1000, IP Action: no, Offload: no
UDP    192.168.203.198:56050 ->192.168.203.194:111 Forward I    1
UDP    192.168.203.194:111 ->192.168.203.198:56050 Forward 0    1

```

user@host> show services sessions application-protocol dns

Interface name: ms-2/0/0

```

Session: 293, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    50.50.50.2:43677 -> 60.60.60.10:53 Forward I      1
UDP    60.60.60.10:53 -> 1.1.1.1:43677 Forward 0      1
Session: 53, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    50.50.50.2:37494 -> 60.60.60.10:53 Forward I      1
UDP    60.60.60.10:53 -> 1.1.1.1:37494 Forward 0      1
Session: 66, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    50.50.50.2:48161 -> 60.60.60.10:53 Forward I      1
UDP    60.60.60.10:53 -> 1.1.1.1:48161 Forward 0      1
Session: 17, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    50.50.50.2:38908 -> 60.60.60.10:53 Forward I      1
UDP    60.60.60.10:53 -> 1.1.1.1:38908 Forward 0      1
Session: 42, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    50.50.50.2:58189 -> 60.60.60.10:53 Forward I      1
UDP    60.60.60.10:53 -> 1.1.1.1:58189 Forward 0      1

```

user@host> show services sessions application-protocol ftp

Interface name: ms-4/1/0

```

Session: 1, ALG: 1, Flags: 0x0040, IP Action: no, Offload: no
TCP    30.1.1.1:32843 -> 20.1.1.1:21 Forward I      26
TCP    20.1.1.1:21 -> 1.1.1.0:32843 Forward 0      30

```

user@host> show services sessions application-protocol pptp

Interface name: ms-2/0/0

```

Session: 3, ALG: pptp, Flags: 0x2800, IP Action: no, Offload: no, Asymmetric: no
GRE    40.40.40.10:0 -> 15.15.15.10:0 Forward 0      21
GRE    15.15.15.10:0 -> 40.40.40.10:65000 Forward I      0
Session: 2, ALG: pptp, Flags: 0x2800, IP Action: no, Offload: no, Asymmetric: no
GRE    15.15.15.10:0 -> 40.40.40.10:49913 Forward I      88
GRE    40.40.40.10:49913 -> 15.15.15.10:65001 Forward 0      0
Session: 1, ALG: pptp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP    15.15.15.10:1511 -> 40.40.40.10:1723 Forward I      13
TCP    40.40.40.10:1723 -> 15.15.15.10:1511 Forward 0      12

```

user@host> show services sessions application-protocol rtsp

Interface name: ms-0/1/0

```

Session: 13, ALG: rtsp, Flags: 0x0800, IP Action: no, Offload: no
UDP    9.1.0.2:5004 -> 9.0.0.2:3989 Forward 0      152
UDP    9.0.0.2:3989 -> 3.1.2.1:5004 Forward I      0
Session: 9, ALG: rtsp, Flags: 0x0800, IP Action: no, Offload: no
UDP    9.1.0.2:5004 -> 9.0.0.2:3986 Forward 0      3
UDP    9.0.0.2:3986 -> 3.1.2.1:5004 Forward I      0

```

user@host> show services sessions application-protocol rsh

```

Interface name: ms-2/0/0
Session: 3, ALG: 2, Flags: 0x0840, IP Action: no, Offload: no
TCP    60.60.60.10:1023 ->    50.50.50.2:1020 Forward 0      4
TCP    50.50.50.2:1020 ->    60.60.60.10:1023 Forward I      3
Session: 1, ALG: 2, Flags: 0x0040, IP Action: no, Offload: no
TCP    50.50.50.2:1021 ->    60.60.60.10:514 Forward I    1331
TCP    60.60.60.10:514 ->    50.50.50.2:1021 Forward 0    2485

user@host> show services sessions application-protocol sip
Interface name: ms-2/0/0
Session: 4, ALG: sip, Flags: 0x0800, IP Action: no, Offload: no
UDP    20.1.1.2:6000 ->    30.1.1.2:12682 Forward I      246
UDP    30.1.1.2:12682 ->    70.1.1.2:6000 Forward 0      0
Session: 1, ALG: sip, Flags: 0x0000, IP Action: no, Offload: no
UDP    20.1.1.2:5060 ->    30.1.1.2:5060 Forward I      10
UDP    30.1.1.2:5060 ->    70.1.1.2:5060 Forward 0      9

user@host> show services sessions application-protocol sql
Interface name: ms-2/0/0
Session: 3934, ALG: sqlnet, Flags: 0x0800, IP Action: no, Offload: no
TCP    50.50.50.2:39754 ->   40.40.40.10:1408 Forward I      26
TCP    40.40.40.10:1408 ->   1.1.1.1:39754 Forward 0      23

user@host> show services sessions application-protocol talk
Interface name: ms-0/2/0
Session: 4, ALG: 65, Flags: 0x0800, IP Action: no, Offload: no
TCP    2.2.2.2:36888 ->    1.1.1.2:33294 Forward 0      4
TCP    1.1.1.2:33294 ->    2.2.2.2:36888 Forward I      3
Session: 7, ALG: 65, Flags: 0x0800, IP Action: no, Offload: no
UDP    2.2.2.2:1165 ->    1.1.1.2:518 Forward 0      1
UDP    1.1.1.2:518 ->    2.2.2.2:1165 Forward I      1
Session: 8, ALG: 65, Flags: 0x0000, IP Action: no, Offload: no
UDP    1.1.1.2:1509 ->    2.2.2.2:518 Forward I      3
UDP    2.2.2.2:518 ->    1.1.1.2:1509 Forward 0      3
Session: 6, ALG: 0, Flags: 0x0000, IP Action: no, Offload: no
UDP    1.1.1.1:123 ->    1.1.1.2:123 Forward 0      4

```

#### show services sessions count

```

user@host> show services sessions count
Interface  Service set      Sessions count
ms-1/1/0   ss                2

```

#### show services sessions destination-port

```

user@host> show services sessions destination-port 21
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP    10.2.2.2:52138 ->    10.1.1.2:21 Forward I      25
TCP    10.1.1.2:21 ->    10.2.2.2:52138 Forward 0      24

```

#### show services sessions destination-prefix

```

user@host> show services sessions destination-prefix 10.1.1.2
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP    10.2.2.2:52138 ->    10.1.1.2:21 Forward I      25
TCP    10.1.1.2:21 ->    10.2.2.2:52138 Forward 0      24

```

#### show services sessions interface

```

user@host> show services sessions interface ms-1/1/0

```

```

ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP      10.2.2.2:52138 ->      10.1.1.2:21    Forward I      30
TCP      10.1.1.2:21    ->      10.2.2.2:52138 Forward 0      29

```

#### show services sessions protocol

```

user@host> show services sessions protocol tcp
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP      10.2.2.2:52138 ->      10.1.1.2:21    Forward I      30
TCP      10.1.1.2:21    ->      10.2.2.2:52138 Forward 0      29

```

#### show services sessions service-set

```

user@host> show services sessions service-set sample
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP      10.2.2.2:52138 ->      10.1.1.2:21    Forward I      33
TCP      10.1.1.2:21    ->      10.2.2.2:52138 Forward 0      31

```

#### show services sessions source-port

```

user@host> show services sessions source-port 21
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP      10.2.2.2:52138 ->      10.1.1.2:21    Forward I      33
TCP      10.1.1.2:21    ->      10.2.2.2:52138 Forward 0      31

```

#### show services sessions source-prefix

```

user@host> show services sessions source-prefix 10.2.2.2
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP      10.2.2.2:52138 ->      10.1.1.2:21    Forward I      33
TCP      10.1.1.2:21    ->      10.2.2.2:52138 Forward 0      31

```

## show services software

<b>Syntax</b>	<b>show services software</b> <b>&lt;count&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 10.4. <count> option added in Junos OS Release 11.2.
<b>Description</b>	Display information about software services. Information is displayed on both 6rd and DS-Lite services.
<b>Options</b>	<b>count</b> <i>interface-name</i> — (Optional) Display the current software counts for a service set for both DS-Lite and 6rd.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services software on page 1210</a> <a href="#">show services software count on page 1210</a>
<b>Output Fields</b>	<a href="#">Table 78 on page 1210</a> lists the output fields for the <b>command-name</b> command. Output fields are listed in the approximate order in which they appear.

**Table 78: show-services-software Output Fields**

Field Name	Field Description	Level of Output
<b>Interface</b>	Interface for which information is displayed.	All levels
<b>Service Set</b>	Service set containing the software rules for the interface.	All levels
<b>Software</b>	Name of the software concentrator.	All levels
<b>Direction</b>	Direction of the flow.	All levels
<b>Flow count</b>	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 1212](#)  
[show services software flows count on page 1212](#)  
[show services software flows ds-lite B4 on page 1213](#)  
[show services software flows ds-lite AFTR on page 1213](#)  
[services software flows ds-lite AFTR and B4 on page 1213](#)

**Output Fields** [Table 79 on page 1212](#) lists the output fields for the `show services software flows` command. Output fields are listed in the approximate order in which they appear.

Table 79: show services software flows Output Fields

Field Name	Field Description
<b>Interface</b>	Name of the interface.
<b>Service set</b>	Name of the service set.
<b>Flow</b>	Description of flow, including protocol input and output interface addresses.
<b>State</b>	Flow state. Value is: <ul style="list-style-type: none"> <li>• <b>Forward</b></li> </ul>
<b>Dir</b>	Flow direction. Values are: <ul style="list-style-type: none"> <li>• <b>I</b>—inbound</li> <li>• <b>O</b>—outbound</li> </ul>
<b>Frm count</b>	Number of frames transferred.
<b>NAT dest</b>	NAT translation of the decapsulated address.
<b>Software</b>	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

<b>Syntax</b>	<pre>show services software statistics &lt;ds-lite&gt; &lt;ds-lite&gt; &lt;interface interface-name&gt; &lt;v6rd&gt;</pre>
<b>Release Information</b>	Command introduced in Junos OS Release 10.4.
<b>Description</b>	Display information about software services.
<b>Options</b>	<p><b>ds-lite</b>—(Optional) Display only DS-Lite.</p> <p><b>interface interface-name</b> —(Optional) Name of the interface servicing the software. When you omit this option, data for all interfaces are shown.</p> <p><b>v6rd</b>—(Optional) Display only 6rd statistics.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services software statistics on page 1217</a> <a href="#">show services software statistics ds-lite on page 1218</a>
<b>Output Fields</b>	Table 80 on page 1214 lists the output fields for the <b>command-name</b> command. Output fields are listed in the approximate order in which they appear.

Table 80: 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 80: command-name Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Slow Path Packets Processed for EIF/HP</b>	Number of slow path EIF/HP packets processed.	<b>statistics for ds-lite only</b>
<b>Fast Path Packets Processed</b>	Number of packets processed that are not <i>slow path</i> .	<b>statistics</b>
<b>Fast Path Encapsulated</b>	Number of packets encapsulated in the fast path.	<b>statistics</b>
<b>Softwire EIF Accept</b>	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.	<b>statistics for ds-lite only</b>
<b>Rule Match Succeeded</b>	Number of packets that matched a softwire rule.	<b>statistics</b>
<b>Rule Match Failed</b>	Number of packets that did not match any softwire rule.	<b>statistics</b>
<b>IPv6 Packets Fragmented</b>	Number of packets fragmented by the services PIC.	<b>statistics for ds-lite only</b>
<b>IPv4 Client Fragments</b>	Number of IPv4 fragments received from the client end over the softwire tunnel destined to the server.	<b>statistics for ds-lite only</b>
<b>IPv4 Server First Fragments</b>	Number of IPv4 first fragments received from the server destined to go over the softwire tunnel to the client.	<b>statistics for ds-lite only</b>
<b>IPv4 Server More Fragments</b>	Number of IPv4 other fragments (excluding first and last fragment) received from the server destined to go over the softwire tunnel to the client.	<b>statistics for ds-lite only</b>
<b>IPv4 Server Last Fragments</b>	Number of IPv4 last fragments received from the server destined to go over the softwire tunnel to the client.	<b>statistics for ds-lite only</b>
<b>ICMPv4 Packets sent</b>	Number of ICMPv4 packets sent to the softwire concentrator.	<b>statistics</b>
<b>ICMPv4 Error Packets sent</b>	Number of ICMPv4 error packets sent to the softwire concentrator.	<b>statistics</b>
<b>ICMPv6 Packets sent</b>	Number of ICMPv6 packets sent to the softwire concentrator.	<b>statistics</b>
<b>Dropped ICMPv6 packets destined to AFTR</b>	Number of ICMPv6 packets dropped instead of sending to the softwire concentrator.	<b>statistics</b>
<b>Softwire Creation Failed</b>	Number of softwire creation failures.	<b>statistics for ds-lite and 6rd</b>

Table 80: command-name Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Softwire Creation Failed for EIF/HP</b>	Number of softwire creation failures for EIF/HP.	<b>statistics for ds-lite only</b>
<b>Flow Creation Failed</b>	Number of flow creation failures.	<b>statistics</b>
<b>Flow Creation Failed for EIF/HP</b>	Number of flow creation failures for EIF/HP.	<b>statistics for ds-lite only</b>
<b>Flow Creation Failed - Retry</b>	Number of flow creations retried after failure.	<b>statistics</b>
<b>Slow Path Failed</b>	Number of failures detected in the slow path.	<b>statistics</b>
<b>Slow Path Failed - Retry</b>	Number of times processing of a packet was reprocessed in the slow path.	<b>statistics</b>
<b>Packet not IPv4-in-IPv6</b>	Number of IPv4 packets not encapsulated in IPv6.	<b>statistics for ds-lite only</b>
<b>IPv6 Fragmentation Error</b>	Number of IPv6 packets with fragmentation errors.	<b>statistics</b>
<b>Slow Path Failed-IPv6 Next Header Offset</b>	Number of IPv6 header errors detected in slow path processing.	<b>statistics for ds-lite only</b>
<b>Decapsulated Packet not IPv4</b>	Number of packets without IPv4 inner header.	<b>statistics for ds-lite only</b>
<b>Decap Failed - IPv6 Next Header Offset</b>	Decapsulation failure due to an unexpected inner header.	<b>statistics for ds-lite only</b>
<b>Decap Failed - IPv4 L3 Integrity</b>	Decapsulation failure due to incorrect Layer 3 data, such as not an IP packet, bad source or destination address, checksum error, or protocol error.	<b>statistics for ds-lite only</b>
<b>Decap Failed - IPv4 L4 Integrity</b>	Decapsulation failure due to incorrect Layer 4 data, such as errors in TCP, UDP, or TCP headers.	<b>statistics for ds-lite only</b>
<b>No Softwire ID</b>	Number of times a softwire ID was not found.	<b>statistics</b>
<b>No Flow Extension</b>	Number of times flow extensions were not found.	<b>statistics</b>
<b>ICMPv4 Dropped Packets</b>	Number of ICMPv4 packets dropped.	<b>statistics</b>
<b>Packet not IPv6-in-IPv4</b>	Number of IPv6 packets not encapsulated in IPv4.	<b>statistics for v6rd only</b>

Table 80: command-name Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Decapsulated Packet not IPv6</b>	Number of packets without an IPv6 inner header.	<b>statistics for v6rd only</b>
<b>Encapsulation Failed - No packet memory</b>	Failed to encapsulate IPv6 packets in IPv4 due to low memory.	<b>statistics for v6rd only</b>
<b>Flow limit exceeded</b>	Flow not created because configured maximum flows per software is exceeded.	<b>statistics</b>
<b>Session limit exceeded</b>	Flow not created because configured maximum DS-Lite software sessions per IPv6 prefix is exceeded.	<b>statistics for ds-lite only</b>

## 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&gt; 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 1223](#)  
[show services stateful-firewall conversations destination-port on page 1223](#)

**Output Fields** [Table 81 on page 1222](#) 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 81: show services stateful-firewall conversations Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of an adaptive services interface.
<b>Service set</b>	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.
<b>Conversation</b>	Information about a group of related flows. <ul style="list-style-type: none"> <li>• <b>ALG Protocol</b>—Application-level gateway protocol.</li> <li>• <b>Number of initiators</b>—Number of flows that initiated a session.</li> <li>• <b>Number of responders</b>—Number of flows that responded in a session.</li> </ul>
<b>Flow or Flow Prot</b>	Protocol used for this flow.
<b>Source</b>	Source prefix of the flow, in the format <i>source-prefix-port</i> .
<b>Destination</b>	Destination prefix of the flow.
<b>State</b>	Status of the flow: <ul style="list-style-type: none"> <li>• <b>Drop</b>—Drop all packets in the flow without response.</li> <li>• <b>Forward</b>—Forward the packet in the flow without looking at it.</li> <li>• <b>Reject</b>—Drop all packets in the flow with response.</li> <li>• <b>Watch</b>—Inspect packets in the flow.</li> </ul>
<b>Dir</b>	Direction of the flow: input (I) or output (O).
<b>Source NAT</b>	Original and translated source IPv4 or IPv6 addresses are displayed if Network Address Translation (NAT) is configured on this particular flow or conversation.
<b>Frm Count</b>	Number of frames in the flow.
<b>Destin NAT</b>	Original and translated destination IPv4 or IPv6 addresses are displayed if NAT is configured on this particular flow or conversation.



Table 81: 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: <b>Yes</b> or <b>No</b> .
TCP window size	Negotiated TCP connection window size, in bytes.
TCP acknowledge	TCP acknowledgment sequence number.
TCP tickle	Whether TCP inquiry mode is on ( <b>enabled</b> or <b>disabled</b> ) 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

<b>Syntax</b>	<code>show services stateful-firewall flow-analysis</code> <code>&lt;interface <i>interface-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.4R1.
<b>Description</b>	Display stateful firewall flow statistics.
<b>Options</b>	<b>none</b> —Display standard information about all stateful firewall flow statistics.  <b>interface <i>interface-name</i></b> —(Optional) Display information about a particular interface. .
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services stateful-firewall flow-analysis on page 1225</a> <a href="#">show services stateful-firewall flow-analysis interface sp-3/0/0 on page 1226</a>
<b>Output Fields</b>	<a href="#">Table 82 on page 1224</a> lists the output fields for the <b>show services stateful-firewall flow-analysis</b> command. Output fields are listed in the approximate order in which they appear.

**Table 82: 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 82: 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
60 - 120      :0            0
30 - 60       :0            0
15 - 30       :0            7330
5 - 15        :0            0
1 - 5         :0            0
0 - 1         :0            7330

```

## 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 1006](#)

**List of Sample Output** [show services stateful-firewall flows on page 1231](#)  
[show services stateful-firewall flows \(For Software Flows\) on page 1231](#)  
[show services stateful-firewall flows brief on page 1232](#)  
[show services stateful-firewall flows extensive on page 1232](#)  
[show services stateful-firewall flows count on page 1232](#)  
[show services stateful-firewall flows destination port on page 1232](#)  
[show services stateful-firewall flows source port on page 1232](#)  
[show services stateful-firewall flows \(Twice NAT\) on page 1232](#)

**Output Fields** [Table 83 on page 1230](#) 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 83: show services stateful-firewall flows Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of the interface.
<b>Service set</b>	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.
<b>Flow Count</b>	Number of flows in a session.
<b>Flow or Flow Prot</b>	Protocol used for this flow.
<b>Source</b>	Source prefix of the flow in the format <i>source-prefix:port</i> . For ICMP flows, port information is not displayed.



Table 83: show services stateful-firewall flows Output Fields (*continued*)

Field Name	Field Description
<b>Dest</b>	Destination prefix of the flow. For ICMP flows, port information is not displayed.
<b>State</b>	Status of the flow: <ul style="list-style-type: none"> <li>• <b>Drop</b>—Drop all packets in the flow without response.</li> <li>• <b>Forward</b>—Forward the packet in the flow without looking at it.</li> <li>• <b>Reject</b>—Drop all packets in the flow with response.</li> <li>• <b>Watch</b>—Inspect packets in the flow.</li> </ul>
<b>Dir</b>	Direction of the flow: input (I) or output (O).
<b>Frm count</b>	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@host> 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

```

### show services stateful-firewall flows source port

```

user@host> 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

```

### show services stateful-firewall flows (Twice NAT)

```

user@host> 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 1008](#)

**List of Sample Output** [show services stateful-firewall sip-call extensive on page 1237](#)

**Output Fields** [Table 84 on page 1236](#) 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 84: show services stateful-firewall sip-call Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of an adaptive services interface.
<b>Service set</b>	Name of a service set.
<b>From</b>	Initiator address.
<b>To</b>	Responder address.
<b>Call ID</b>	SIP call identification string.
<b>Number of initiator flows</b>	Number of <b>control</b> , <b>contact</b> , or <b>media</b> initiator flows.
<b>Number of responder flows</b>	Number of <b>control</b> , <b>contact</b> , or <b>media</b> responder flows.
<b>protocol</b>	Protocol used for this flow.
<b>source-prefix</b>	Source prefix of the flow in the format <b>source-prefix : port</b> .
<b>destination-prefix</b>	Destination prefix of the flow.
<b>state</b>	Status of the flow: <ul style="list-style-type: none"> <li>• <b>Drop</b>—Drop all packets in the flow without a response.</li> <li>• <b>Forward</b>—Forward the packet in the flow without examining it.</li> <li>• <b>Reject</b>—Drop all packets in the flow with a response.</li> <li>• <b>Unknown</b>—Unknown status.</li> <li>• <b>Watch</b>—Inspect packets in the flow.</li> </ul>
<b>direction</b>	Direction of the flow: input (I), output (O), or unknown (U).

Table 84: show services stateful-firewall sip-call Output Fields (*continued*)

Field Name	Field Description
<i>frame-count</i>	Number of frames in the flow.
<b>Byte count</b>	Number of bytes forwarded in the flow.
<b>Flow role</b>	Role of the flow that is under evaluation: <b>Initiator</b> , <b>Master</b> , <b>Responder</b> , or <b>Unknown</b> .
<b>Timeout</b>	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 1011](#)

**List of Sample Output** [show services stateful-firewall sip-register extensive on page 1241](#)

**Output Fields** [Table 85 on page 1241](#) 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 85: show services stateful-firewall sip-register Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of an adaptive services interface.
<b>Service set</b>	Name of a service set.
<b>SIP Register</b>	Register information header.
<b>Protocol</b>	Protocol used for this flow.
<b>Registered IP</b>	Register IP address.
<b>Port</b>	Register port number.
<b>Expiration timeout</b>	Configured lifetime, in seconds.
<b>Timeout remaining</b>	Lifetime remaining, in seconds.
<b>From</b>	Initiator address.
<b>To</b>	Responder address.
<b>Call ID</b>	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

<b>Syntax</b>	<pre>show services stateful-firewall statistics &lt;application-protocol <i>protocol</i>&gt; &lt;brief   detail   extensive   summary&gt; &lt;interface <i>interface-name</i>&gt; &lt;service-set <i>service-set</i>&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display stateful firewall statistics.
<b>Options</b>	<p><b>none</b>—Display standard information about all stateful firewall statistics.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(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><b>service-set <i>service-set</i></b>—(Optional) Display information about a particular service set.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear services stateful-firewall statistics on page 1014</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show services stateful-firewall statistics extensive on page 1250</a>
<b>Output Fields</b>	Table 86 on page 1243 lists the output fields for the <b>show services stateful-firewall statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 86: show services stateful-firewall statistics Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of an adaptive services interface.
<b>Service set</b>	Name of a service set.
<b>New flows</b>	Rule match counters for new flows: <ul style="list-style-type: none"> <li><b>Rule Accepts</b>—New flows accepted.</li> <li><b>Rule Discards</b>—New flows discarded.</li> <li><b>Rule Rejects</b>—New flows rejected.</li> </ul>
<b>Existing flow types packet counters</b>	Rule match counters for existing flows: <ul style="list-style-type: none"> <li><b>Accepts</b>—Match existing forward or watch flow.</li> <li><b>Drop</b>—Match existing discard flow.</li> <li><b>Rejects</b>—Match existing reject flow.</li> </ul>

Table 86: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
<b>Hairpinning Counters</b>	<p>Hairpinning counters:</p> <ul style="list-style-type: none"> <li>• <b>Slow Path Hairpinned Packets</b>—Slow path packets that were hairpinned back to the internal network.</li> <li>• <b>Fast Path Hairpinned Packets</b>—Fast path packets that were hairpinned back to the internal network.</li> </ul>
<b>Drops</b>	<p>Drop counters:</p> <ul style="list-style-type: none"> <li>• <b>IP option</b>—Packets dropped in IP options processing.</li> <li>• <b>TCP SYN defense</b>—Packets dropped by SYN defender.</li> <li>• <b>NAT ports exhausted</b>—Hide mode. The router has no available Network Address Translation (NAT) ports for a given address or pool.</li> <li>• <b>Sessions dropped due to subscriber flow limit</b>—Sessions dropped because the subscriber's flow limit was exceeded.</li> </ul>
<b>Errors</b>	<p>Total errors, categorized by protocol:</p> <ul style="list-style-type: none"> <li>• <b>IP</b>—Total IP version 4 errors.</li> <li>• <b>TCP</b>—Total Transmission Control Protocol (TCP) errors.</li> <li>• <b>UDP</b>—Total User Datagram Protocol (UDP) errors.</li> <li>• <b>ICMP</b>—Total Internet Control Message Protocol (ICMP) errors.</li> <li>• <b>Non-IP packets</b>—Total non-IPv4 errors.</li> <li>• <b>ALG</b>—Total application-level gateway (ALG) errors</li> </ul>

Table 86: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
IP Errors	<p>IPv4 errors:</p> <ul style="list-style-type: none"> <li>• <b>IP packet length inconsistencies</b>—IP packet length does not match the Layer 2 reported length.</li> <li>• <b>Minimum IP header length check failures</b>—Minimum IP header length is 20 bytes. The received packet contains less than 20 bytes.</li> <li>• <b>Reassembled packet exceeds maximum IP length</b>—After fragment reassembly, the reassembled IP packet length exceeds 65,535.</li> <li>• <b>Illegal source address 0</b>—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.</li> <li>• <b>Illegal destination address 0</b>—Destination address is not a valid address. The address is reserved.</li> <li>• <b>TTL zero errors</b>—Received packet had a time-to-live (TTL) value of 0.</li> <li>• <b>Illegal IP protocol number (0 or 255)</b>—IP protocol is 0 or 255.</li> <li>• <b>Land attack</b>—IP source address is the same as the destination address.</li> <li>• <b>Non-IPv4 packets</b>—Packet was not IPv4. (Only IPv4 is supported.)</li> <li>• <b>Bad checksum</b>—Packet had an invalid IP checksum.</li> <li>• <b>Illegal IP fragment length</b>—Illegal fragment length. All fragments (other than the last fragment) must have a length that is a multiple of 8 bytes.</li> <li>• <b>IP fragment overlap</b>—Fragments have overlapping fragment offsets.</li> <li>• <b>IP fragment reassembly timeout</b>—Some of the fragments for an IP packet were not received in time, and the reassembly handler dropped partial fragments.</li> <li>• <b>IP fragment limit exceeded: 0</b>—Fragments that exceeded the limit.</li> <li>• <b>Unknown: 0</b>—Unknown fragments.</li> </ul>

Table 86: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
TCP Errors	



Table 86: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
	TCP protocol errors:
	<ul style="list-style-type: none"> <li>• <b>TCP header length inconsistencies</b>—Minimum TCP header length is 20 bytes, and the IP packet received does not contain at least 20 bytes.</li> <li>• <b>Source or destination port number is zero</b>—TCP source or destination port is zero.</li> <li>• <b>Illegal sequence number and flags combinations</b> — Dropped because of TCP errors, such as an illegal sequence number, which causes an illogical combination of flags to be set.</li> <li>• <b>SYN attack (multiple SYN messages seen for the same flow)</b>—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.</li> <li>• <b>First packet not a SYN message</b>—First packets for a connection are not SYN packets. These packets might originate from previous connections or from someone performing an ACK/FIN scan.</li> <li>• <b>TCP port scan (TCP handshake, RST seen from server for SYN)</b>—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).</li> <li>• <b>Bad SYN cookie response</b>—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.</li> <li>• <b>TCP reconstructor sequence number error</b>—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.</li> <li>• <b>TCP reconstructor retransmissions</b>—This counter is incremented for the retransmitted packets during connection 3-way handshake.</li> <li>• <b>TCP partially opened connection timeout (SYN)</b>—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.</li> <li>• <b>TCP partially opened connection timeout (SYN-ACK)</b>—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.</li> <li>• <b>TCP partially closed connection reuse</b>—Not supported.</li> <li>• <b>TCP 3-way error - client sent SYN+ACK</b>—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.</li> <li>• <b>TCP 3-way error - server sent ACK</b>—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.</li> <li>• <b>TCP 3-way error - SYN seq number retransmission mismatch</b>—This counter is incremented when the SYN is received again with a different sequence number from the first SYN sequence number.</li> <li>• <b>TCP 3-way error - RST seq number mismatch</b>—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</li> </ul>

Table 86: 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"> <li>• <b>TCP 3-way error - FIN received</b>—This counter is incremented when the FIN is received during the 3-way handshake.</li> <li>• <b>TCP 3-way error - invalid flags (PSH, URG, ECE, CWR)</b>—This counter is incremented when any of the PSH, URG, ECE, or CWR flags were received during the 3-way handshake.</li> <li>• <b>TCP 3-way error - SYN recvd but no client flows</b>—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.</li> <li>• <b>TCP 3-way error - first packet SYN+ACK</b>—The first packet received was SYN+ACK instead of SYN.</li> <li>• <b>TCP 3-way error - first packet FIN+ACK</b>—The first packet received was FIN+ACK instead of SYN.</li> <li>• <b>TCP 3-way error - first packet FIN</b>—The first packet received was FIN instead of SYN.</li> <li>• <b>TCP 3-way error - first packet RST</b>—The first packet received was RST instead of SYN.</li> <li>• <b>TCP 3-way error - first packet ACK</b>—The first packet received was ACK instead of SYN.</li> <li>• <b>TCP 3-way error - first packet invalid flags (PSH, URG, ECE, CWR)</b>—The first packet received had invalid flags.</li> <li>• <b>TCP Close error - no final ACK</b>—This counter is incremented when ACK is not received after the FINs are received from both directions.</li> <li>• <b>TCP Resumed Flow</b>—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.</li> </ul>
UDP Errors	<p>UDP protocol errors:</p> <ul style="list-style-type: none"> <li>• <b>IP data length less than minimum UDP header length (8 bytes)</b>—Minimum UDP header length is 8 bytes. The received IP packets contain less than 8 bytes.</li> <li>• <b>Source or destination port is zero</b>—UDP source or destination port is 0.</li> <li>• <b>UDP port scan (ICMP error seen for UDP flow)</b>—ICMP error is received for a UDP flow. This could be a genuine UDP flow, but it is counted as an error.</li> </ul>
ICMP Errors	<p>ICMP protocol errors:</p> <ul style="list-style-type: none"> <li>• <b>IP data length less than minimum ICMP header length (8 bytes)</b>—ICMP header length is 8 bytes. This counter is incremented when received IP packets contain less than 8 bytes.</li> <li>• <b>ICMP error length inconsistencies</b>—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.</li> <li>• <b>Duplicate ping sequence number</b>—Received ping packet has a duplicate sequence number.</li> <li>• <b>Mismatched ping sequence number</b>—Received ping packet has a mismatched sequence number.</li> <li>• <b>No matching flow</b>—No matching existing flow was found for the ICMP error.</li> </ul>

Table 86: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
<b>ALG errors</b>	<p>Accumulation of all the application-level gateway protocol (ALG) drops counted separately in the ALG context:</p> <ul style="list-style-type: none"> <li>• <b>BOOTP</b>—Bootstrap protocol errors</li> <li>• <b>DCE-RPC</b>—Distributed Computing Environment-Remote Procedure Call protocols errors</li> <li>• <b>DCE-RPC portmap</b>—Distributed Computing Environment-Remote Procedure Call protocols portmap service errors</li> <li>• <b>DNS</b>—Domain Name System protocol errors</li> <li>• <b>Exec</b>—Exec errors</li> <li>• <b>FTP</b>—File Transfer Protocol errors</li> <li>• <b>H323</b>—H.323 standards errors</li> <li>• <b>ICMP</b>—Internet Control Message Protocol errors</li> <li>• <b>IIOB</b>—Internet Inter-ORB Protocol errors</li> <li>• <b>Login</b>—Login errors</li> <li>• <b>NetBIOS</b>—NetBIOS errors</li> <li>• <b>Netshow</b>—NetShow errors</li> <li>• <b>Real Audio</b>—RealAudio errors</li> <li>• <b>RPC</b>—Remote Procedure Call protocol errors</li> <li>• <b>RPC portmap</b>—Remote Procedure Call protocol portmap service errors</li> <li>• <b>RTSP</b>—Real-Time Streaming Protocol errors</li> <li>• <b>Shell</b>—Shell errors</li> <li>• <b>SIP</b>—Session Initiation Protocol errors</li> <li>• <b>SNMP</b>—Simple Network Management Protocol errors</li> <li>• <b>SQLNet</b>—SQLNet errors</li> <li>• <b>TFTP</b>—Trivial File Transfer Protocol errors</li> <li>• <b>Traceroute</b>—Traceroute errors</li> </ul>
<b>Drop Flows</b>	<ul style="list-style-type: none"> <li>• <b>Maximum Ingress Drop flows allowed</b>—Maximum number of ingress flow drops allowed.</li> <li>• <b>Maximum Egress Drop flows allowed</b>—Maximum number of egress flow drops allowed.</li> <li>• <b>Current Ingress Drop flows</b>—Current number of ingress flow drops.</li> <li>• <b>Current Egress Drop flows</b>—Current number of egress flow drops.</li> <li>• <b>Ingress Drop Flow limit drops count</b>—Number of ingress flow drops due to maximum number of ingress flow drops being exceeded.</li> <li>• <b>Egress Drop Flow limit drops count</b>—Number of egress flow drops due to maximum number of egress flow drops being exceeded.</li> </ul>

## 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
Hairpinning 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

<b>Syntax</b>	show services stateful-firewall application-protocol sip
<b>Release Information</b>	Command introduced in Junos OS Release 7.4.
<b>Description</b>	Display stateful firewall Session Initiation Protocol (SIP) statistics.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services stateful-firewall statistics application-protocol-sip on page 1253</a>
<b>Output Fields</b>	<a href="#">Table 87 on page 1252</a> lists the output fields for the <b>show services stateful-firewall statistics application-protocol-sip</b> command. Output fields are listed in the approximate order in which they appear.

**Table 87: 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 87: show services stateful-firewall statistics application-protocol-sip**  
**Output Fields (continued)**

Field Name	Field Description
<b>OPTIONS</b>	Number of new, invalid, and retransmitted requests to query SIP capabilities.
<b>INFO</b>	Number of new, invalid, and retransmitted requests carrying application-level information.
<b>UPDATE</b>	Number of new, invalid, and retransmitted SIP dialogue updates.
<b>REFER</b>	Number of new, invalid, and retransmitted requests to the recipient to contact a third party.
<b>Provisional responses</b>	Number of new, invalid, and retransmitted responses from the user agent server to indicate the progress of a SIP transaction.
<b>OK responses to INVITES</b>	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.
<b>OK responses to non-INVITES</b>	OK responses to SIP messages other than an Invite message.
<b>Redirection responses</b>	Responses from the user agent server to a user agent client requesting the client to contact a different SIP uniform resource identifier (URI).
<b>Request failure responses</b>	Responses that indicate a definite failure from a particular server. The client must not retry the same request without modification after receiving this response.
<b>Server failure responses</b>	Responses that indicate a server failure.
<b>Global failure responses</b>	Responses that indicate a server has definitive information about a particular user, not just the particular instance indicated in the Request URI.
<b>Invalid responses</b>	Responses that are invalid.
<b>Response (all) retransmits</b>	Retransmissions of all responses.
<b>Parser</b>	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		
ACK	0	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

<b>Syntax</b>	show services stateful-firewall subscriber analysis <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 11.4.
<b>Description</b>	Display information about the number of active subscribers on the service physical interface card (PIC).
<b>Options</b>	<b>none</b> —Display standard information about all active subscribers on the PIC.  <b>interface <i>interface-name</i></b> —(Optional) Display information about a particular interface.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show services stateful-firewall subscriber analysis on page 1256</a> <a href="#">show services stateful-firewall subscriber-analysis on page 1256</a>
<b>Output Fields</b>	<a href="#">Table 88 on page 1255</a> lists the output fields for the <b>show services stateful-firewall subscriber analysis</b> command. Output fields are listed in the approximate order in which they appear.

**Table 88: 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



## PART 12

# Index

- [Index on page 1261](#)



# Index

## Symbols

#, comments in configuration statements.....	xxxvi
( ), in syntax descriptions.....	xxxvi
6rd flows	
statistics.....	1211
< >, in syntax descriptions.....	xxxvi
[ ], in configuration statements.....	xxxvi
{ }, in configuration statements.....	xxxvi
(pipe), in syntax descriptions.....	xxxvi

## A

adaptive services interfaces.....	1029
status information, displaying.....	1029
adaptive-services-pics statement.....	709
address statement	
interfaces	
usage guidelines.....	24
NAT.....	710
voice services.....	710
usage guidelines.....	666
address-allocation statement.....	711
address-range statement	
NAT.....	711
Aggregated Multiservices interfaces	
load balancing.....	1077
aggregation statement.....	712
usage guidelines.....	387
alert (system logging severity level).....	26, 683
ALGs	
configuring.....	326
supported on the MS-MIC and MS-MPC.....	350
ALGsJ	
default.....	127, 321
allow-ip-options statement.....	713
usage guidelines.....	362
allow-multicast statement.....	714
usage guidelines.....	17
allow-overlapping-nat-pools statement.....	714
AMS	
HA.....	649, 651
NAT.....	649, 660

AMS interfaces	
service sets for dynamic NAT, resplit the NAT	
pool.....	13
anomaly checklist.....	356
anti-replay-window-size statement.....	714, 715
usage guidelines.....	458, 463
any (system logging severity level).....	26, 682
any-any match condition	
Ipsec.....	454
app-mapping-timeout statement.....	716
application statement.....	717
usage guidelines.....	325
application-profile statement.....	720
usage guidelines.....	559
application-protocol statement.....	718
usage guidelines.....	326
application-set statement.....	721
usage guidelines.....	325
application-sets statement	
CoS.....	721
IDS.....	722
usage guidelines.....	386
NAT.....	722
usage guidelines.....	56
stateful firewall.....	723
usage guidelines.....	360
usage guidelines.....	557
applications.....	386
example configuration.....	343
applications statement	
application-level gateways.....	725
applications hierarchy.....	723
CoS.....	724
IDS.....	724
usage guidelines.....	386
NAT.....	725
usage guidelines.....	56
stateful firewall.....	725
usage guidelines.....	360
usage guidelines.....	557
applying service set to interface.....	9
AS PIC	
multicast traffic.....	17
redundancy.....	20, 686
associations, clearing.....	974
authentication statement.....	726
usage guidelines.....	418

authentication-algorithm statement	
IKE.....	727
usage guidelines.....	436
IPsec.....	728
usage guidelines.....	446
authentication-method statement.....	729
usage guidelines.....	436
auxiliary-spi statement.....	730
usage guidelines.....	418
<b>B</b>	
backup AS PIC.....	20
backup Link Services IQ PIC.....	592
backup-remote-gateway statement.....	730
usage guidelines.....	457
bandwidth	
and delay buffer allocation.....	575
guaranteed.....	575, 580
basic-nat-pt option	
configuring.....	129
basic-nat44 option	
configuring.....	79
basic-nat66 option	
configuring.....	85
braces, in configuration statements.....	xxxvi
brackets	
angle, in syntax descriptions.....	xxxvi
square, in configuration statements.....	xxxvi
bundle statement.....	731
usage guidelines.....	670
by-destination statement.....	731
usage guidelines.....	387
by-pair statement.....	732
usage guidelines.....	387
by-source statement.....	733
usage guidelines.....	387
bypass-traffic-on-exceeding-flow-limits	
statement.....	733
bypass-traffic-on-pic-failure statement.....	734
usage guidelines.....	9
<b>C</b>	
certificates	
for IKE negotiation, displaying.....	1117
PKI	
CA certificates, clearing.....	995
CA certificates, displaying.....	1082
CA certificates, loading manually.....	1017
certificate revocation lists, clearing.....	997
certificate revocation lists,	
displaying.....	1088
certificate revocation lists, loading	
manually.....	1019
key pair, generating.....	1022
local certificates, clearing.....	998, 999
local certificates, displaying.....	1090
local certificates, loading manually.....	1026
local certificates, requesting	
manually.....	1020, 1025
local certificates, requesting online.....	1016
local certificates, requesting that CA	
install.....	1023
local certificates, requests, clearing.....	996
local certificates, requests,	
displaying.....	1086
cg-nat statement.....	734
NAT.....	61
CGNAT	
ALGs.....	127, 321
CIR.....	580
cisco-interoperability statement.....	735
usage guidelines.....	591
class statement.....	736
clear security pki ca-certificate command.....	995
clear security pki certificate-request	
command.....	996
clear security pki crl command.....	997
clear security pki key-pair.....	998
clear security pki local-certificate command.....	999
clear services cos statistics command.....	963
clear services crtp statistics command.....	964
clear services ids command.....	965
clear services ids destination-table	
command.....	966
clear services ids pair-table command.....	967
clear services ids source-table command.....	968
clear services inline nat pool command.....	969
clear services inline nat statistics command.....	970
clear services inline software statistics	
command.....	971
clear services ipsec-vpn certificates	
command.....	972
clear services ipsec-vpn ike security-associations	
command.....	973
clear services ipsec-vpn ipsec security-associations	
command.....	974
clear services ipsec-vpn ipsec statistics	
command.....	975



clear services l2tp destination command.....	976
clear services l2tp destination statistics command.....	977
clear services l2tp multilink command.....	978
clear services l2tp session command.....	979
clear services l2tp session statistics command.....	981
clear services l2tp tunnel command.....	983
clear services l2tp tunnel statistics command.....	985
clear services nat flows command.....	987
clear services nat mappings app command.....	990, 991, 993
clear services nat mappings command.....	988
clear services service-sets statistics packet-drops command.....	1001
clear services service-sets statistics syslog command.....	1002
clear services sessions command.....	1003
clear services stateful-firewall flows command.....	1006
clear services stateful-firewall sip-call command.....	1008
clear services stateful-firewall sip-register command.....	1011
clear services stateful-firewall statistics command.....	1014
clear-dont-fragment-bit (NAT option).....	738
clear-dont-fragment-bit statement for NAT options.....	738
GRE tunnel.....	737
IPsec.....	737
usage guidelines.....	456
service-set.....	738
usage guidelines.....	457, 464
clear-ike-sas-on-pic-restart statement.....	739
usage guidelines.....	420
clear-ipsec-sas-on-pic-restart statement.....	739
usage guidelines.....	420
comments, in configuration statements.....	xxxvi
compression statement.....	740
usage guidelines.....	667, 668
compression-device statement.....	740
usage guidelines.....	670
configuring NAT-PT with DNS application-level gateways example.....	136
conventions text and syntax.....	xxxv
cookies, SYN.....	383
copy-dont-fragment-bit statement IPsec.....	741
service-set.....	741
CoS action statements.....	558
applications.....	557
example configuration.....	560
link services interfaces.....	570, 606
match conditions.....	557
rules.....	561
CoS services clear statistics.....	963
mapping, displaying code point aliases to bit patterns.....	1093
critical (system logging severity level).....	26, 683
CRTP services flows, displaying.....	1098
output, displaying.....	1096
statistics, clearing.....	964
curly braces, in configuration statements.....	xxxvi
customer support.....	xxxvii
contacting JTAC.....	xxxvii
<b>D</b> data statement.....	742
usage guidelines.....	559
dead peer detection (DPD) protocol.....	457
delay buffer calculating.....	575, 580
shaping rate.....	575, 580
delay-buffer-rate statement usage guidelines.....	575
description statement IKE.....	743
usage guidelines.....	443
IPsec.....	743
usage guidelines.....	448, 450
destination NAT configuring.....	97, 173, 176
destination-address statement CoS.....	743
IDS.....	744
usage guidelines.....	386
IPsec.....	744
usage guidelines.....	454
NAT.....	745
usage guidelines.....	56

stateful firewall.....	745	DS-Lite flows	
usage guidelines.....	360	statistics.....	1211
usage guidelines.....	557	ds-lite statement.....	760
destination-address-range statement		usage guidelines.....	237
IDS.....	746	dscp statement.....	761
usage guidelines.....	386	usage guidelines.....	558
NAT.....	747	dynamic address-only source translation	
usage guidelines.....	56	configuring.....	181
stateful firewall.....	748	dynamic authentication.....	496
usage guidelines.....	360	dynamic NAT	
destination-pool statement.....	748	configuring.....	181
usage guidelines.....	57	dynamic route insertion.....	497
destination-port range statement		dynamic rules.....	496
NAT.....	750	dynamic security associations	
destination-port statement		usage guidelines.....	420, 435
applications.....	723	dynamic statement.....	761
RPM.....	749	usage guidelines.....	420
usage guidelines.....	331	dynamic-nat44 option	
destination-prefix statement.....	750, 751	usage guidelines.....	181
usage guidelines.....	387		
destination-prefix-ipv6 statement.....	751	<b>E</b>	
usage guidelines.....	387	ecmp-alb statement.....	762
destination-prefix-list statement		ei-mapping-timeout statement.....	763
CoS.....	752	emergency (system logging severity	
IDS.....	752	level).....	26, 682
NAT.....	753	enable-change-on-ams-redistribution	
stateful firewall.....	753	statement.....	764
usage guidelines.....	360	usage guidelines.....	13
destined-port statement		enable-rejoin statement	
NAT.....	754	aggregated Multiservices.....	765
deterministic-port-block-allocation		encapsulation statement.....	766
statement.....	755	voice services	
dh-group statement.....	756	usage guidelines.....	669
usage guidelines.....	437	encryption statement.....	767
dial-options statement.....	757	usage guidelines.....	419
interfaces		encryption-algorithm statement	
usage guidelines.....	684	IKE.....	768
direction statement.....	758	usage guidelines.....	438
usage guidelines.....	416	IPsec.....	768
dnat-44 option		usage guidelines.....	448
usage guidelines.....	97, 173, 176	error (system logging severity level).....	27, 683
documentation		establish-tunnels statement.....	769
comments on.....	xxxvii	event policy	
drop-member-traffic statement		all (tracing flag).....	29
aggregated Multiservices.....	759	configuration (tracing flag).....	29
ds-lite		database (tracing flag).....	29
subnet session limitation		events (tracing flag).....	29
configuring.....	253	policy (tracing flag).....	29

**F**

f-max-period statement.....	769
usage guidelines.....	667
facility-override statement.....	770, 771
usage guidelines.....	26
family statement	
aggregated Multiservices.....	771
interfaces.....	772
usage guidelines.....	24
voice services.....	773
filters	
used with services.....	9
firewall filters	
service filters.....	18
floods	
SYN.....	383
flow collector services	
statistics	
dropped-packet, clearing.....	1001, 1002
flow limiting.....	15
font conventions.....	xxxv
force-entry statement.....	774
usage guidelines.....	387
forwarding classes	
fragmentation.....	570
forwarding classes, displaying.....	1093
forwarding-class statement.....	774, 775
usage guidelines.....	558, 570
fragment-threshold statement	
LSQ.....	776
usage guidelines.....	570
voice services.....	777
usage guidelines.....	668
fragmentation	
forwarding classes.....	570
multiclass MLPPP.....	606
fragmentation and reassembly.....	668
fragmentation-map statement.....	777
usage guidelines.....	570
fragmentation-maps statement.....	778
usage guidelines.....	570
FRF.12.....	668
example configuration.....	629
LSQ.....	626
FRF.16.....	615
configuration example.....	618
from statement	
CoS.....	779
HCM.....	781

IDS.....	780
usage guidelines.....	384, 386
IPsec.....	781
usage guidelines.....	452, 454
NAT.....	782
usage guidelines.....	56
stateful firewall.....	783
usage guidelines.....	359, 360
usage guidelines.....	556
ftp statement.....	784
usage guidelines.....	559

**G**

guaranteed rate.....	580
guaranteed-rate statement	
usage guidelines.....	580

**H**

hash-keys statement	
aggregated multiservices.....	785
hello-interval statement	
L2TP.....	789
usage guidelines.....	681
hide-avps statement.....	790
usage guidelines.....	682
high-availability-options statement	
aggregated Multiservices.....	791
hint statement.....	792
host statement.....	793
HCM.....	794
L2TP.....	792
usage guidelines.....	26, 682
hot-standby statement.....	794

**I**

ICMP	
ALGs, supported on the MS-MIC and	
MS-MPC.....	350
icmp-code statement.....	795
usage guidelines.....	329
icmp-type statement.....	795
usage guidelines.....	329
IDS	
action statements.....	387
applications.....	386
example configurations.....	392
match conditions.....	386
rules.....	384

IDS events		IKE security associations	
clearing		clearing.....	420
for a destination.....	966	ike statement.....	797
for interfaces and services.....	965	usage guidelines.....	435
for source addresses.....	968	ike-access-profile statement.....	798
for source and destination pairs.....	967	usage guidelines.....	462, 499
displaying.....	1101	inactivity-timeout statement.....	798
ids-rule-sets statement		usage guidelines.....	334
usage guidelines.....	14	info (system logging severity level).....	27, 683
ids-rules statement.....	796	initiate-dead-peer-detection statement.....	799
usage guidelines.....	14	usage guidelines.....	458
ignore-entry statement.....	774	inline LSQ services.....	607
usage guidelines.....	387	inline MLPPP for WAN interfaces.....	603
IKE.....	402, 435	inline NAT	
adaptive services interfaces		statistics, displaying.....	1109, 1111
security associations, clearing.....	973	inline software	
security associations, displaying.....	1120	statistics, displaying.....	1114
statistics, clearing.....	975	input statement	
authentication algorithm		interfaces.....	799
usage guidelines.....	436	usage guidelines.....	9, 17
authentication-method statement		inside and outside interfaces.....	12
usage guidelines.....	436	inside-service-interface statement	
DH (Diffie-Hellman) group		usage guidelines.....	13
usage guidelines.....	437	interchassis LSQ failover.....	589
dynamic SAs.....	435	interface preservation.....	595
encryption-algorithm statement		interface statement	
usage guidelines.....	438	service interface pool.....	800
lifetime		interface style service sets.....	12
usage guidelines.....	438	interface-service statement.....	800
mode statement		usage guidelines.....	9
usage guidelines.....	441	interfaces statement	
policy.....	439	aggregated Multiservices.....	801
example.....	444	voice services.....	802
policy statement		Internet Key Exchange See IKE	
usage guidelines.....	439	intrachassis LSQ failover.....	592
pre-shared-key statement		intrusion detection	
usage guidelines.....	441	example configurations.....	392
proposals statement		rule set.....	392
usage guidelines.....	441	invalid SPI recovery	
supported software standards.....	408	enabling.....	444
version statement		IPsec	
usage guidelines.....	441	action statements.....	456
IKE profile		authentication statement	
configuring access profile.....	497	usage guidelines.....	418
IKE proposal		authentication-algorithm statement	
example configuration.....	439	usage guidelines.....	446
IKE proposals		direction	
default.....	500	usage guidelines.....	416
		dynamic authentication.....	496

dynamic endpoints for IPsec tunnels.....	495	IPsec services	
dynamic endpoints interface		adaptive services interfaces	
configuration.....	499	backup and primary, switching	
dynamic rules.....	496	tunnels.....	1028
dynamic security associations		IKE security associations, clearing.....	973
usage guidelines.....	420	IKE security associations, displaying.....	1120
encryption		IPSec security associations, clearing.....	974
usage guidelines.....	419	IPSec security associations,	
encryption-algorithm statement		displaying.....	1124
usage guidelines.....	448	IPSec statistics, clearing.....	975
example policy configuration.....	452	IPSec statistics, displaying.....	1128
IKE.....	402	ipsec statement.....	803
lifetime of SA.....	448	usage guidelines.....	445
lifetime-seconds statement.....	448	ipsec-inside-interface	
match conditions.....	454	usage guidelines.....	496
minimum configurations		ipsec-inside-interface statement.....	803
dynamic SA .....	414	usage guidelines.....	454
manual SA .....	413	ipsec-interface-id statement	
overview.....	401	usage guidelines.....	499
perfect-forward-secrecy statement		ipsec-vpn-options statement.....	804
usage guidelines.....	451	usage guidelines.....	461
policy		ipsec-vpn-rule-sets statement	
overview.....	450	usage guidelines.....	14
policy statement		ipsec-vpn-rules statement.....	804
usage guidelines.....	450	usage guidelines.....	14
proposal statement		IPv4	
usage guidelines.....	445	napt-44 option.....	113
proposals statement		translation type	
usage guidelines.....	451	basic-nat-pt option.....	129
protocol statement (dynamic SA)		basic-nat44 option.....	79
usage guidelines.....	449	basic-nat66 option.....	85
protocol statement (manual SA)		IPv4 dynamic source translation	
usage guidelines.....	417	configuring.....	113
rule sets.....	459	IPv4 static source translation	
security associations.....	402	AMS.....	660
security parameter index		example.....	660
usage guidelines.....	417	IPv6	
service set dynamic endpoints		napt-66 option.....	117
configuration.....	499	IPv6 dynamic source translation	
Services SDK		configuring.....	117
configuration.....	549	ipv6-multicast-interfaces statement.....	805
supported software standards.....	408		
IPsec proposals		<b>J</b>	
default.....	500	Junos Network Secure.....	355
IPsec rules		overview.....	355
match directions.....	454	See also stateful firewall	

**L****L2TP**

access profile.....	679, 680
attribute-value pairs.....	682
example configuration.....	686
redundancy.....	686
timers.....	681

**L2TP LAC services**

destination	
clearing.....	976, 977

**L2TP services**

multilink sessions	
clearing.....	978
displaying.....	1135
RADIUS information.....	1139
session statistics	
clearing.....	981
sessions	
clearing.....	979
displaying.....	1143
summary information, displaying.....	1151
tunnel statistics, clearing.....	985
tunnels, clearing.....	983
tunnels, displaying.....	1156
user information, displaying.....	1162

**L2TP statements**

LAC	
traceoptions.....	935

**LNS**

l2tp-access-profile.....	805
local-gateway.....	811
service-interface.....	889
traceoptions.....	935

l2tp-access-profile statement.....	805
usage guidelines.....	680

**l2tp-interface-id statement**

usage guidelines.....	684
-----------------------	-----

**l2tp-profile statement**

usage guidelines.....	679
-----------------------	-----

**learn-sip-register statement.....**

LFI.....	621, 626, 668
example configuration.....	624, 629

**lifetime-seconds statement**

IKE.....	807
usage guidelines.....	438
IPsec.....	807
usage guidelines.....	448

**limiting flows per service set.....****link PIC redundancy.....****link services interfaces**

CoS components.....	570, 606
---------------------	----------

**link services IQ interfaces.....**

example configuration.....	612, 618
link state replication.....	595
link-layer overhead.....	605
status information, displaying.....	1037

**link state replication**

LSQ PICs.....	595
---------------	-----

**link-layer overhead**

link services IQ interfaces.....	605
----------------------------------	-----

**link-layer-overhead statement.....**

usage guidelines.....	568, 572, 605
-----------------------	---------------

**load balancing.....****load-balance statement.....****load-balancing-options statement**

aggregated Multiservices.....	809
-------------------------------	-----

**local-certificate statement.....**

usage guidelines.....	442
-----------------------	-----

**local-gateway address statement**

usage guidelines.....	680
-----------------------	-----

**local-gateway statement.....**

usage guidelines.....	461
-----------------------	-----

**local-id statement.....**

usage guidelines.....	443
-----------------------	-----

**log output**

adaptive services.....	29
------------------------	----

**log-prefix statement.....**

L2TP.....	812
-----------	-----

usage guidelines.....	26, 682
-----------------------	---------

**logging statement.....**

usage guidelines.....	387
-----------------------	-----

**logical interface scheduling.....****LSQ**

CPU usage information, displaying.....	1131
--	------

**LSQ bandwidth**

oversubscribing.....	575
----------------------	-----

**LSQ failover**

interchassis.....	589
-------------------	-----

stateful intrachassis.....	592
----------------------------	-----

stateless intrachassis.....	592
-----------------------------	-----

**LSQ PICs.....**

redundancy.....	592
-----------------	-----

**lsq-failure-options statement.....**

usage guidelines.....	590
-----------------------	-----

**M****manual security association.....**

- manual statement.....815
  - usage guidelines.....415
- manuals
  - comments on.....xxxvii
- many-to-one statement
  - aggregated Multiservices.....816
- mapping-timeout statement.....818
- match direction usage in service sets.....12
- match-direction statement
  - CoS.....818
  - IDS.....819
    - usage guidelines.....384
  - IPsec.....819
    - usage guidelines.....452
  - NAT.....820
  - stateful firewall.....820
    - usage guidelines.....360
  - usage guidelines.....557
- max-drop-flows statement.....821
- max-flows statement.....822
  - usage guidelines.....15
- max-sessions-per-subscriber statement.....806, 823
- maximum-contexts statement.....824
  - usage guidelines.....667
- maximum-send-window statement.....824
  - usage guidelines.....681
- member-failure-options statement
  - aggregated Multiservices.....825
- member-interface statement
  - aggregated Multiservices.....827
- mlfr-uni-nni-bundles-inline.....829
- MLPPP.....609, 621
  - configuration example.....612
  - example configuration.....624
- mode statement.....830
  - usage guidelines.....441
- MS-MIC and MS-MPC
  - supported ALGs
    - identical support as uKernel.....350
- MS-MPC
  - configuration example
    - napt.....120
- mss statement.....830
  - usage guidelines.....387
- multi-link-layer-2-inline.....831
- multicast traffic
  - AS PIC.....17
- multiclass MLPPP
  - fragmentation.....606
- multilink bundles
  - fractional T1.....621
    - example configuration.....624, 626, 629
  - FRF.12.....626
    - example configuration.....629
  - MLPPP.....621
    - example configuration.....624
  - NxT1.....609, 615
    - configuration example.....612, 618
- multilink-class statement.....831
  - usage guidelines.....606
- multilink-max-classes statement.....832
  - usage guidelines.....606
- MultiServices PIC
  - hardware requirements.....3
- N**
- NAPT
  - comparison of implementation methods.....112
  - configuring.....113, 117
  - IPv4.....113
  - IPv6.....117
  - port allocation
    - round-robin.....104
    - sequential.....104
  - port block allocation.....107
- napt
  - configuration example.....120
- napt-44 option
  - usage guidelines.....113
- napt-66 option
  - usage guidelines.....117
- napt-pt option
  - example.....136
- NAT
  - action statements.....57
  - ALGs.....127, 321
  - AMS.....649
  - applications.....56
  - destination NAT.....97, 173
  - dynamic address-only source translation.....181
  - dynamic NAT.....181
  - dynamic source translation.....113, 117
  - inline.....189
    - configuring.....191
  - inter-chassis high availability.....262
  - ipv6-multicast-interfaces information,
    - displaying.....1169
    - load balancing, example.....660

mapping information, address-pooling	
paired.....	1171
mapping information,	
displaying.....	1166, 1168, 1171
mapping information, endpoint-independent	
.....	1171
mapping information, pcp.....	1171
match conditions.....	56
NAPT	
configuring address pools.....	103
NAT-PT example.....	136
overview.....	43
service sets.....	61
session logging.....	219
static destination address translation.....	97, 173
status information, displaying.....	1176
twice NAT	
description.....	37, 46
nat	
flows	
clearing.....	987
mappings	
clearing.....	988, 990, 991, 993
nat-options statement.....	832
nat-rule-sets statement	
usage guidelines.....	14
nat-rules statement.....	833
usage guidelines.....	14
network address translation	
configuration example	
napt.....	120
network address translation See NAT	
next-hop style service sets.....	13
next-hop-service statement.....	834
usage guidelines.....	11
no-anti-replay statement.....	835
usage guidelines.....	458, 463
no-fragmentation statement.....	836
usage guidelines.....	570
no-ipsec-tunnel-in-traceroute statement.....	836
usage guidelines.....	466
no-per-unit-scheduler statement.....	837
no-termination-request statement.....	837
usage guidelines.....	590
no-translation statement.....	838
usage guidelines.....	57
notice (system logging severity level).....	27, 683
NxTI bundles	
FRF.16.....	615
configuration example.....	618
MLPPP.....	609
configuration example.....	612
<b>O</b>	
output statement.....	838
usage guidelines.....	9, 17
outside-service-interface statement	
usage guidelines.....	13
overload-pool statement.....	839
usage guidelines.....	57
overload-prefix statement.....	839
usage guidelines.....	57
oversubscription.....	575
<b>P</b>	
packet-based IPsec.....	454
parentheses, in syntax descriptions.....	xxxvi
passive-mode-tunneling statement.....	840
usage guidelines.....	465
per-unit scheduling.....	842
per-unit-scheduler statement.....	842
usage guidelines.....	575, 580, 609, 615
perfect-forward-secrecy statement.....	843
usage guidelines.....	451
pgcp statement	
NAT.....	844
pgcp-rules statement	
service-set.....	844
ping	
ALGs, supported on the MS-MIC and	
MS-MPC.....	350
PIR.....	575
PKI See certificates, PKI	
policy statement	
IKE.....	845
usage guidelines.....	439
IPsec.....	846
usage guidelines.....	450
pool statement.....	847
service interface pool.....	848
port block allocation.....	107
deterministic.....	108
algorithms.....	108
configuring.....	171
interim syslog messages.....	219



- secured.....107, 163
    - configuring.....169
  - secured, guidelines for configuring.....164, 167
  - Port Control Protocol
    - Configuring.....153
    - Configuring a Service Set to Apply PCP.....155
    - Configuring PCP Server Options.....153, 155
  - port forwarding
    - configuring.....176
    - dnat-44.....173
    - static destination address translation.....173
    - without destination address translation.....176
  - port forwarding without static destination address translation
    - configuring.....176
  - port statement
    - NAT.....849
    - voice services.....851
    - usage guidelines.....667
  - port-forwarding
    - example.....177
  - port-forwarding statement
    - destined-port statement.....754
    - NAT.....852
    - translated-port statement.....942
  - port-forwarding-mappings statement.....852
  - ports-per-session statement.....853
  - post-service-filter statement.....853
    - usage guidelines.....9
  - ppp-access-profile statement.....854
    - usage guidelines.....680
  - ppp-profile statement
    - usage guidelines.....679
  - pre-shared-key statement.....854
    - usage guidelines.....441
  - preserve-interface statement.....855
    - usage guidelines.....595
  - primary statement
    - link services.....856
      - usage guidelines.....593
    - services PIC.....855
      - usage guidelines.....20
  - proposal statement
    - IKE.....856
      - usage guidelines.....435
    - IPsec.....857
      - usage guidelines.....445
  - proposals statement
    - IKE.....857
      - usage guidelines.....441
    - IPsec.....857
      - usage guidelines.....451
  - protocol statement
    - applications.....858
      - usage guidelines.....328
    - IPsec.....859
      - usage guidelines.....417, 449
  - ptsp-rule-sets statement
    - usage guidelines.....15
  - ptsp-rules statement.....859
    - usage guidelines.....15
- ## Q
- queues statement.....860
    - usage guidelines.....667
- ## R
- RADIUS information
    - displaying.....1139
  - RADIUS servers
    - configuration example.....660
  - random-allocation statement.....849
  - reassembly-timeout statement.....860
  - receive-window statement.....861
    - usage guidelines.....681
  - redistribute-all-traffic statement
    - aggregated Multiservices.....861
  - redundancy
    - AS PIC.....20
    - L2TP.....686
  - redundancy-options statement.....862, 863
    - usage guidelines.....20
  - redundant adaptive services interfaces
    - reverting to the primary interface.....1015
    - status information, displaying.....1061
    - switching to the secondary interface.....1015
  - redundant link services IQ interfaces
    - status information, displaying.....1063
  - reflexive | reverse statement.....864
    - usage guidelines.....560
  - rejoin-timeout statement
    - aggregated Multiservices.....865
  - remote-gateway statement.....865
    - usage guidelines.....457
  - remote-id statement.....866
    - usage guidelines.....443

remotely-controlled statement.....	866	rule-set statement.....	879
request interface (revert   switchover) (Adaptive Services) command.....	1015	CoS	
request security pki ca-certificate enroll command.....	1016	usage guidelines.....	561
request security pki ca-certificate load command.....	1017	IDS.....	880
request security pki ca-certificate verify command.....	1018	usage guidelines.....	392
request security pki crt load command.....	1019	IPsec.....	880
request security pki generate-certificate-request command.....	1020	usage guidelines.....	459
request security pki generate-key-pair command.....	1022	NAT.....	881, 951
request security pki local-certificate enroll command.....	1023	software.....	882
request security pki local-certificate generate-self-signed command.....	1025	stateful firewall.....	881
request security pki local-certificate load command.....	1026	usage guidelines.....	363
request security pki local-certificate verify command.....	1027		
request services ipsec-vpn ipsec switch tunnel command.....	1028	<b>S</b>	
request-url statement.....	867	secondary statement	
respond-bad-ip statement		link services.....	883
usage guidelines.....	444	usage guidelines.....	593
respond-bad-spi statement		services PIC.....	882
IKE.....	869	usage guidelines.....	20
retransmit-interval statement.....	869	secure-nat-mapping statement.....	883
usage guidelines.....	681	secured-port-block-allocation statement.....	884
rpc-program-number statement.....	870	security associations	
usage guidelines.....	342	clearing.....	420
rtp statement.....	871	configuring.....	415
usage guidelines.....	667	server (PCP) statement.....	886
rule statement		service filters.....	18
CoS.....	872	service interface configuration.....	9
IDS.....	873	service rules configuration.....	14
usage guidelines.....	384	service sets	
IPsec.....	875	example configuration.....	20
usage guidelines.....	452	overview.....	4
NAT.....	877	service statement.....	887
software.....	229, 879	usage guidelines.....	17
stateful firewall.....	878	service-domain statement.....	888
usage guidelines.....	359	usage guidelines.....	11
usage guidelines.....	556	service-filter statement.....	888
		firewall	
		usage guidelines.....	18
		interfaces	
		usage guidelines.....	9
		service-interface statement.....	889
		usage guidelines.....	9, 680
		service-interface-pools statement.....	890
		service-set options	
		bouncing or reset of service sets for AMS	
		interfaces.....	13
		service-set statement.....	890, 891
		NAT.....	61
		usage guidelines.....	17

service-set statements	
Adaptive Services interfaces	
service-interface.....	889
services sets	
CPU usage, displaying.....	1184
dropped packet statistics	
clearing.....	1001
displaying.....	1192
memory usage, displaying.....	1186
statistics of packets dropped owing to integrity errors in protocol headers	
displaying.....	1188
summary information, displaying.....	1200
syslog statistics	
clearing.....	1002
displaying.....	1194
services statement	
APPID	
usage guidelines.....	704
L2TP	
usage guidelines.....	682
NAT.....	893
service sets	
usage guidelines.....	26
session logging.....	219
session-limit statement.....	894
usage guidelines.....	387
set-dont-fragment-bit statement	
IPsec.....	895
service-set.....	895
shaping-rate statement	
usage guidelines.....	568, 575, 580
show interfaces (Adaptive Services)	
command.....	1029
show interfaces (Link Services IQ) command.....	1037
show interfaces (Redundant Adaptive Services)	
command.....	1061
show interfaces (Redundant Link Services IQ)	
command.....	1063
show interfaces command.....	1077
show interfaces redundancy command.....	1080
show security pki ca-certificate command.....	1082
show security pki certificate-request	
command.....	1086
show security pki crl command.....	1088
show security pki local-certificate	
command.....	1090
show services cos statistics command.....	1093
show services crtp command.....	1096
show services crtp flows command.....	1098
show services hcm statistics rule command.....	1100
show services ids command.....	1101
show services inline nat pool command.....	1109
show services inline nat statistics command.....	1111
show services inline softwire statistics	
command.....	1114
show services ipsec-vpn certificates	
command.....	1117
show services ipsec-vpn ike security-associations	
command.....	1120
show services ipsec-vpn ipsec security-associations	
command.....	1124
show services ipsec-vpn ipsec statistics	
command.....	1128
show services l2tp multilink command.....	1135
show services l2tp radius command.....	1139
show services l2tp session command.....	1143
show services l2tp summary command.....	1151
show services l2tp tunnel command.....	1156
show services l2tp user command.....	1162
show services link-services cpu-usage	
command.....	1131
show services nat deterministic-nat internal-host	
command.....	1166
show services nat deterministic-nat nat-port-block	
command.....	1168
show services nat ipv6-multicast-interfaces	
command.....	1169
show services nat mappings command.....	1171
show services nat pool command.....	1176
show services pcp statistics command.....	1181
show services service-sets cpu-usage	
command.....	1184
show services service-sets memory-usage	
command.....	1186
show services service-sets statistics integrity-drops	
command.....	1188
show services service-sets statistics packet-drops	
command.....	1192
show services service-sets statistics syslog	
command.....	1194
show services service-sets statistics tcp-mss	
command.....	1199
show services service-sets summary	
command.....	1200
show services sessions command.....	1202
show services softwire command.....	1210
show services softwire flows command.....	1211

show services softwire statistics command.....	1214	source-port statement	
show services stateful-firewall conversations		RPM.....	906
command.....	1220	usage guidelines.....	331
show services stateful-firewall flow-analysis		source-prefix statement.....	906, 907
command.....	1224	usage guidelines.....	387
show services stateful-firewall flows		source-prefix-ipv6 statement.....	907
command.....	1228	usage guidelines.....	387
show services stateful-firewall sip-call		source-prefix-list statement	
command.....	1234	CoS.....	908
show services stateful-firewall sip-register		IDS.....	908
command.....	1239	NAT.....	909
show services stateful-firewall statistics		stateful firewall.....	909
application-protocol sip command.....	1252	usage guidelines.....	360
show services stateful-firewall statistics		sip statement.....	910
command.....	1243	usage guidelines.....	417
show services stateful-firewall subscriber analysis		stateful firewall	
command.....	1255	action statements.....	362
sip statement.....	897	anomalies.....	356
usage guidelines.....	559	applications.....	360
sip-call-hold-timeout statement.....	896	conversations	
snmp-command statement.....	897	displaying.....	1220
usage guidelines.....	342	example configuration.....	363
snmp-trap-thresholds statement.....	898	flow analysis	
softwire flows		displaying.....	1224
statistics.....	1211	flows	
softwire-concentrator statement.....	899	clearing.....	1006
softwire-rules statement.....	900	displaying.....	1228
usage guidelines.....	15	match conditions.....	360
source-address statement		overview.....	355
CoS.....	901	rules.....	363
IDS.....	902	SIP call information	
usage guidelines.....	386	clearing.....	1008
IPsec.....	902	displaying.....	1234
usage guidelines.....	454	SIP register information	
NAT.....	903	clearing.....	1011
usage guidelines.....	56	displaying.....	1239
service-set system log.....	901	SIP statistics	
stateful firewall.....	903	displaying.....	1252
usage guidelines.....	360	statistics	
usage guidelines.....	557	clearing.....	1014
source-address-range statement		displaying.....	1243
IDS.....	904	subscriber analysis	
usage guidelines.....	386	displaying.....	1255
NAT.....	904	stateful NAT64	
usage guidelines.....	56	configuring.....	75
stateful firewall.....	905	stateful-firewall-rule-sets statement	
usage guidelines.....	360	usage guidelines.....	15
source-pool statement.....	905	stateful-firewall-rules statement.....	910
usage guidelines.....	57	usage guidelines.....	15

- stateful-nat64 statement.....911
  - statement
    - IPsec
      - usage guidelines.....466
    - L2TP
      - usage guidelines.....689
  - static destination address translation
    - configuring.....97, 173
  - subnet session limitation
    - ds-lite
      - configuring.....253
  - support, technical See technical support
  - SYN cookies.....383
  - SYN floods
    - SYN cookies.....383
  - syn-cookie statement.....916
    - usage guidelines.....387
  - syntax conventions.....xxxv
  - syslog statement.....911
    - IDS.....912
      - usage guidelines.....387
    - IPsec.....912
      - usage guidelines.....456, 459
    - L2TP.....913
      - usage guidelines.....682
    - NAT.....913
    - service sets.....914
      - usage guidelines.....26
    - stateful firewall.....915
      - usage guidelines.....362
    - usage guidelines.....558
  - system log statement
    - NAT
      - usage guidelines.....57
- T**
- tcp-mss
    - statistics, displaying.....1199
  - tcp-mss statement.....917
  - technical support
    - contacting JTAC.....xxxvii
  - term statement
    - CoS.....918
    - HCM.....922
    - IDS.....919
      - usage guidelines.....384
    - IPsec.....921
      - usage guidelines.....452
    - NAT.....923
    - stateful firewall.....924
      - usage guidelines.....359
    - usage guidelines.....556
  - then statement.....925
    - HCM.....925
    - IDS.....926
      - usage guidelines.....384
    - IPsec.....927
      - usage guidelines.....452
    - NAT.....928
    - stateful firewall.....929
      - usage guidelines.....359, 362
    - usage guidelines.....556
  - threshold statement.....930
    - usage guidelines.....387
  - time-to-live threshold.....342
  - trace-options
    - server (tracing flag).....29
    - timer-events (tracing flag).....29
  - traceoptions statement
    - IPsec.....933
    - L2TP.....935
    - security.....931
    - services.....939
  - traceroute ALG
    - on the MS-MIC and MS-MPC.....350
  - tracing flags
    - event policy
      - all.....29
      - configuration.....29
      - database.....29
      - events.....29
      - policy.....29
      - server.....29
      - timer-events.....29
  - tracing operations
    - adaptive services.....27
  - traffic-control-profiles statement
    - usage guidelines.....575, 580
  - translated statement.....941
    - usage guidelines.....57
  - translated-port statement
    - NAT.....942
  - translation-type statement.....943
    - basic-nat-pt option.....129
    - basic-nat44 option.....79
    - basic-nat66 option.....85
    - dnat-44 option, configuring.....97, 173
    - dynamic-nat44, configuring.....181

napt-44 option, configuring.....	113	voice statement.....	957, 958
napt-66 option, configuring.....	117	usage guidelines.....	559
napt-pt option, example.....	136	<b>W</b>	
stateful-nat64 option, configuring.....	75	warm standby	
usage guidelines.....	57	AS PIC.....	20
transport statement		LSQ PIC.....	592
NAT.....	941	warm-standby statement.....	958
trigger-link-failure statement.....	942	warning (system logging severity level).....	27, 683
usage guidelines.....	590		
trusted-ca statement.....	945		
usage guidelines.....	463		
ttl-threshold statement.....	945		
usage guidelines.....	342		
tunnel-group statement.....	946		
usage guidelines.....	679		
tunnel-mtu statement.....	947, 948		
usage guidelines.....	459, 466		
tunnel-timeout statement.....	949		
usage guidelines.....	681		
twice NAT.....	37, 46		
twice-napt-44 option			
example.....	177		
<b>U</b>			
unit statement			
aggregated Multiservices.....	951		
interfaces.....	952		
link services.....	953		
Universal Unique Identifier.....	342		
url-rule statement.....	950		
url_identifier, statement.....	949, 950		
uuid statement.....	954		
usage guidelines.....	342		
<b>V</b>			
v6rd statement.....	955		
usage guidelines.....	255		
version statement			
IKE.....	956		
usage guidelines.....	441		
video statement.....	956, 957		
usage guidelines.....	559		
voice services			
bundles.....	670		
encapsulation.....	669		
example configuration.....	671		
interface type.....	666		
voice services interfaces			
interleave fragments.....	668		